NOTE

This Microfiche meets the quality requirements of the American National Standard for Microfiche of Documents (ANSI PH5.9-1975). However, because of imperfections in the document from which it was made, it contains images that are illegible or difficult to read. The source document for this microfiche is the best copy that could be obtained prior to filming.

Congressional Information Service, Inc
Washington, D.C.

CIS

SUPERCOMPUTERS

HEARINGS

BEFORE THE

COMMITTEE ON
SCIENCE AND TECHNOLOGY
U.S. HOUSE OF REPRESENTATIVES

NINETY-EIGHTH CONGRESS

FIRST SESSION

NOVEMBER 15, 16, 1983

[No. 47]

Printed for the use of the
Committee on Science and Technology

U.S. GOVERNMENT PRINTING OFFICE
28-293 O WASHINGTON : 1984

H701·20

COMMITTEE ON SCIENCE AND TECHNOLOGY

DON FUQUA, Florida, *Chairman*

ROBERT A. ROE, New Jersey
GEORGE E. BROWN, Jr., California
JAMES H. SCHEUER, New York
RICHARD L. OTTINGER, New York
TOM HARKIN, Iowa
MARILYN LLOYD, Tennessee
DOUG WALGREN, Pennsylvania
DAN GLICKMAN, Kansas
ALBERT GORE, Jr., Tennessee
ROBERT A. YOUNG, Missouri
HAROLD L. VOLKMER, Missouri
BILL NELSON, Florida
STAN LUNDINE, New York
RALPH M. HALL, Texas
DAVE McCURDY, Oklahoma
MERVYN M. DYMALLY, California
PAUL SIMON, Illinois
NORMAN Y. MINETA, California
RICHARD J. DURBIN, Illinois
MICHAEL A. ANDREWS, Texas
BUDDY MacKAY, Florida
TIM VALENTINE, North Carolina
HARRY M. REID, Nevada
ROBERT G. TORRICELLI, New Jersey
FREDERICK C. BOUCHER, Virginia

LARRY WINN, Jr., Kansas
MANUEL LUJAN, Jr., New Mexico
ROBERT S. WALKER, Pennsylvania
WILLIAM CARNEY, New York
F. JAMES SENSENBRENNER, Jr., Wisconsin
JUDD GREGG, New Hampshire
RAYMOND J. McGRATH, New York
JOE SKEEN, New Mexico
CLAUDINE SCHNEIDER, Rhode Island
BILL LOWERY, California
ROD CHANDLER, Washington
HERBERT H. BATEMAN, Virginia
SHERWOOD L. BOEHLERT, New York
ALFRED A. McCANDLESS, California
TOM LEWIS, Florida

J. H. Poore, *Executive Director*
Robert C. Ketcham, *General Counsel*
Regina A. Davis, *Administrator*
David S. Jeffery, *Minority Staff Director*

CONTENTS

WITNESSES

November 15, 1983:

Page
Dr. George Kozmetsky, director, Institute for Constructive Capitalism, University of Texas—Austin 6
Hon. John T. Myers, a Representative in Congress from the State of Indiana 109
Dr. Steven Beering, president, Purdue University, representing: National Association of State Universities and Land-Grant Colleges and Association of American Universities 109

November 16, 1983:

Dr. Edward Knapp, director, National Science Foundation 142
Dr. Edith Martin, Deputy Under Secretary, Research and Advanced Technology, Department of Defense 157

APPENDIX

Additional Testimony Submitted for the Record

Dr. John Brademas, president, New York University 205
Dr. James Browne, professor of computer science, University of Texas-Austin.. 210
Dr. John Gibbons, director, Office of Technology Assessment 217
Donna L. Valtri, project director, Office of Technology Assessment 222
Dr. Paul B. Schneck, division leader, information sciences, Office of Naval Research, representing: Institute of Electrical and Electronics Engineers 234
Dr. Alvy Ray Smith, computer graphics project leader, Lucasfilm, Ltd 241
Dr. George Keyworth, Science Adviser to the President, and Director, Office of Science and Technology Policy, Executive Office of the President 247

Followup questions and answers:

Association of American Universities and National Association of State Universities and Land-Grant Colleges 258
National Science Foundation 266
Department of Defense 281

Additional Material Submitted for the Record

Report of the panel on Large Scale Computing in Science and Engineering, Introduction, Dr. Peter Lax, chairman, December 26, 1982 291
A National Computing Environment for Academic Research, National Science Foundation, July 1983 303
The U.S. Activities Board, IEEE, Scientific Supercomputer Committee report, October 25, 1983 359
Report to the Federal Coordinating Council on Science, Engineering and Technology by the Supercomputer Panel on Recommended Government Actions To Provide Access to Supercomputers 403
Report to the Federal Coordinating Council on Science, Engineering and Technology by the Supercomputer Panel on Recommended Government Actions to Retain U.S. Leadership in Supercomputers 428
Academic supercomputer R&D prospects, David Kuck, October 19, 1983 451
Resource material to Committee on Science and Technology hearings on subjects relating to university supercomputing, Robert W. Hayman, October 21, 1983 456

SUPERCOMPUTERS

TUESDAY, NOVEMBER 15, 1983

House of Representatives,
Committee on Science and Technology,
Washington, D.C.

The committee met, pursuant to call, at 1:30 p.m., in room 2318, Rayburn House Office Building, Hon. Don Fuqua (chairman of the committee) presiding.

Present: Representatives Fuqua, MacKay, Reid, Winn, and McGrath.

Mr. Fuqua. The committee will be in order.

Without objection, permission is granted for television, radio and photographic devices.

The hearings this week will build on the joint hearings in June by the Science and Technology Committee's Subcommittee on Energy Development and Applications and the Subcommittee on Energy Research and Production on "Computers and Their Role in Energy Research: Current Status and Future Needs."

For about a year now, various groups in the Government and learned societies have been studying the advanced computer programs and opportunities presented by the Japanese initiatives and, more importantly, our own desire to solve important societal problems and even our desire to be No. 1 in the computer industry. So today and tomorrow we will hear from some of these groups and what they have decided about how we should proceed. We have statements from others who could not accommodate our changed schedule, and I would like at this time to ask unanimous consent that those statements be made a part of the record.[1]

[Statements begin on page 205.]

We'll also hear from representatives from the private sector and academic communities concerning their views about the current and future needs for supercomputer access and the role of the Federal Government in stimulating expanded use and development of advanced computer technology.

The committee's concerns are illustrated by the following true story. A chemistry researcher down in Alabama is supported by a Federal grant. Even though the grantor knew the research required access to a supercomputer, no funds were provided to pay for the access to one. So the chemist succeeded in getting a separate grant, this one a travel grant to go to West Germany to use a supercomputer readily available to German researchers and their associates—and it also included the Alabama chemist.

---

[1] Written testimony submitted for the record by: Dr. James C. Browne, professor of computer science, University of Texas-Austin; Dr. John H. Gibbons, director, Office of Technology Assessment; Dr. Paul B. Schneck, Division Leader, Information Sciences, Office of Naval Research, representing Institute of Electrical and Electronics Engineers; Dr. Alvy Ray Smith, computer graphics project leader, computer division, Lucasfilm Ltd.

Meanwhile, equally adequate facilities are sitting underutilized at Colorado State, the University of Minnesota, and Purdue, where these schools are making substantial financial sacrifices in an effort to get these tools into American universities. To date, because of myths, catch-22 policies, and lack of focus, Federal funding agencies have done little to help universities either extend their resources or acquire access to existing resources.

Problems important to our society are going unresolved—problems like predicting severe storms or developing new drugs—partly because of severely limited access to advanced computer technology by our scientists who have good ideas about how to solve these problems.

Other important problems are going unresolved because even our most advanced computers are too limited. Simulation of an aircraft design will require a computer 200 times more powerful than any available today.

Yet, even with our commitment to solve such problems, computer developers are very hesitant to go forward because the individual financial risks are large and the collaborators which would make the risks acceptable could be violating antitrust laws.

At the same time, Japan is threatening our entire computer industry by targeting the supercomputer segment.

Following these hearings, the committee will be formulating their own views on programs and initiatives in this area. Thus, we will be especially interested to learn from our witnesses their recommendations and viewpoints concerning four points: One, the solution to problems important to society, including those requiring computers 200 times more powerful than currently available; two, the most effective approaches for enhancing access to advanced computers by university faculty and students; three, the mechanisms for accelerating advanced computer research and development; and four, the role of the Federal Government, industry, and academia in exploring advanced computer technology for the benefit of science, engineering, and society.

And at this time, I would like to recognize the ranking minority member of the committee, Mr. Winn, for any comments that he wishes to make.

Mr. WINN. Thank you, Mr. Chairman, for this opportunity to provide a brief statement—which I will include in the record, and not go over the whole thing—and to welcome our witnesses today.

Mr. Chairman, you are again to be commended for directing the committee's attention to a very timely subject—in this case, supercomputers. And I look forward to hearing from our witnesses today.

Thank you, Mr. Chairman.

[The prepared statement of Mr. Winn follows:]

STATEMENT BY THE HONORABLE LARRY WINN, JR.
HEARING ON SUPERCOMPUTERS
NOVEMBER 15, 1983

THANK YOU, MR. CHAIRMAN, FOR THIS OPPORTUNITY TO PROVIDE A BRIEF STATEMENT AND TO WELCOME OUR WITNESSES TODAY. MR. CHAIRMAN, YOU ARE AGAIN TO BE COMMENDED FOR DIRECTING THE COMMITTEE'S ATTENTION TO A VERY TIMELY SUBJECT, IN THIS CASE, SUPERCOMPUTERS.

FROM THE BEGINNING OF THE COMPUTER AGE, THE UNITED STATES HAS BEEN THE ONLY NATION TO DEVELOP AND SELL SUPERCOMPUTERS. THIS HAS PROVIDED US THE CAPABILITY OF PERFORMING CALCULATIONS AT LEAST FOUR TO FIVE TIMES FASTER THAN THE MOST POWERFUL CONVENTIONAL COMPUTER SYSTEMS. HOWEVER, THIS HAS LED TO TWO DISCERNABLE PROBLEMS. ONE, THE USE OF THE SUPERCOMPUTER IN THE UNITED STATES IS NOT NECESSARILY WIDELY DISTRIBUTED. ACCESS TO SUPERCOMPUTER FACILITIES IS LIMITED PRIMARILY TO WEAPONS RESEARCH WITH ITS NECESSARY SECURITY, TO ONE OR TWO FUNDAMENTAL RESEARCH PROGRAMS WITHIN THE DEPARTMENT OF ENERGY, TO ENVIRONMENTAL FORECASTING, AND TO PROPRIETARY USAGE IN PETROLEUM EXPLORATION. SO A LARGE NUMBER OF BASIC, FUNDAMENTAL RESEARCH PROJECTS PROGRESS AT INEFFICIENT RATES DUE TO LACK OF ADEQUATE COMPUTATIONAL CAPABILITY. I STRONGLY FEEL THAT WE MUST EXPLORE THE AVENUES BY WHICH WE MAKE THIS NEEDED COMPUTER CAPABILITY

AVAILABLE TO OUR SCIENTISTS IN INDUSTRY, COLLEGES AND UNIVERSITIES. IT IS HOPED THAT FEDERAL SUPPORT WILL CONSIST PRIMARILY OF COORDINATION AMONG THOSE WHO CAN JUSTIFY FACILITIES OR ACCESS TO THEM.

THE SECOND AREA OF CONCERN IS THAT THE UNITED STATES LEADERSHIP IN SUPERCOMPUTER DEVELOPMENT AND MARKETING IS BEING CHALLENGED BY FOREIGN COUNTRIES, PARTICULARLY THE JAPANESE. THAT COUNTRY HAS ESTABLISHED A JOINT GOVERNMENT-INDUSTRY PROGRAM WITH A COMMITMENT THAT EXTENDS INTO THE NEXT DECADE. THE PURPOSE OF THE EFFORT IS CLEAR CUT AND WELL-DEFINED. THIS STRONG PROGRAM DEFINITION AND COMMITMENT OF NEARLY ONE-HALF BILLION DOLLARS EXCEEDS ANY COORDINATED PROGRAM IN THIS COUNTRY AT THIS TIME.

I AM CONTINUALLY AMAZED AT THE PROGRESS WHICH THE COMPUTER INDUSTRY HAS MADE IN THE PAST 40 YEARS. THROUGH THE ENCOURAGEMENT OF THE FEDERAL GOVERNMENT AS A PRIME CUSTOMER, THE EXPERTISE OF THE SCIENCE COMMUNITY, THE ENGINEERING SKILLS IN THE UNITED STATES, AND THE FREE ENTERPRISE SYSTEM, WE HAVE WITNESSED FANTASTIC GROWTH. YOU CAN NOW BUY A HOME COMPUTER IN A TOY STORE AND CARRY OUT, UNDER ONE ARM, THE COMPUTER CAPABILITY THAT WAS DELIVERED BY TRUCK IN LARGE PACKAGES TO INDUSTRY TEN YEARS AGO. WE NOW TALK OF MILLIONS OF INSTRUCTIONS BEING EXECUTED PER SECOND AND OF GOALS EXCEEDING THAT BY THREE ORDERS OF MAGNITUDE. THE PRICE AND SIZE OF COMPONENTS HAS BEEN REDUCED DRASTICALLY. DATA PROCESSING CHIPS HAVE INVADED OUR LIVES TO THE EXTENT OF HELPING US OPERATE OUR CARS, PLAY GAMES, OPERATE HOME APPLIANCES, AND

OTHER MORE AWE-INSPIRING OCCURRENCES TOO NUMEROUS TO MENTION.

IT APPEARS THAT ONCE AGAIN, AMERICAN INDUSTRY IS RISING TO THE CHALLENGE. WITHIN THE PAST 24 MONTHS, WE HAVE BEEN ADVISED OF THE FORMATION OF MULTI-COMPANY GROUPS SUCH AS THE MICROELECTRONICS AND COMPUTER TECHNOLOGY CORPORATION, A RESEARCH GROUP CONTAINING OVER ONE DOZEN U.S. COMPANIES; ETA SYSTEMS, INC., A NEW COMPANY SPONSORED BY ONE OF OUR LARGE SUPERCOMPUTER MANUFACTURERS; AND THE SEMICONDUCTOR RESEARCH CORPORATION, A NOT-FOR-PROFIT COMPANY FOUNDED BY A HANDFUL OF COMPUTER CHIP MAKERS. THESE THREE COMPANIES WILL ALL ENGAGE IN BASIC RESEARCH AND WILL STRENGTHEN THE AMERICAN POSITION OF SUPERCOMPUTER SUPERIORITY. IT IS MY HOPE THAT THE CONTINUING AND ANTICIPATED NEEDS OF THE FEDERAL GOVERNMENT WILL BE CHANNELED TO THESE AND OTHER RESEARCH ENTITIES, PROVIDING GUIDANCE ON THE DIRECTIONS IN WHICH THEIR WORK CAN BE PERFORMED.

AGAIN, MR. CHAIRMAN, I DO WANT TO EXPRESS MY APPRECIATION TO YOU FOR HOLDING THESE HEARINGS. I AM CERTAIN THEY WILL BE BENEFICIAL TO US ALL.

Mr. **Fuqua**. Thank you very much, Mr. Winn.

And to lead off, our first witness today, I would like to call on Dr. George Kozmetsky from the Institute of Constructive Capitalism at the University of Texas in Austin.

I might say that Dr. Kozmetsky has had a very varied background in many fields of endeavor, and has spent a considerable amount of time on this particular issue, and is very eminently qualified to be our leadoff witness.

George, we're happy to have you, and pleased to hear from you at this time.

Mr. **Winn**. Mr. Chairman, before the doctor starts, I wonder if he could tell us a little bit about the background of his company. The title is kind of intriguing, and maybe an explanation by the doctor would help us understand exactly what type of work you are in.

STATEMENT OF DR. GEORGE KOZMETSKY, DIRECTOR, INSTITUTE FOR CONSTRUCTIVE CAPITALISM, UNIVERSITY OF TEXAS-AUSTIN

Dr. **Kozmetsky**. The institute——

Mr. **Winn**. Would you pull the mike right up in front of your mouth.

Dr. **Kozmetsky**. The institute——

Mr. **Winn**. Thank you.

Dr. **Kozmetsky** [continuing]. That you referred to is strictly academic, pure research, on the basic system that this country was founded on. It is not interested in propaganda. It is simply interested in doing, first, studying new economic systems. Many of us have realized that the Keynesian—neo-Keynesian—theory has not worked, so we're off on new theories. Second, we are extremely interested in using technology as a resource, and we do a great deal of interuniversity meetings on developing entrepreneurs, especially in high technology, for this Nation.

Mr. **Winn**. Thank you very much.

I am sorry to interrupt, Mr. Chairman, but I just thought maybe an explanation of the—of the company might help us understand the direction that he might take.

And thank you, Doctor.

Dr. **Kozmetsky**. It's—it's a part of the University of Texas at Austin.

Mr. **Fuqua**. And it is part of the University of Texas.

Dr. **Kozmetsky**. Yes.

Mr. Chairman, members of the committee, I'm honored to have this opportunity to appear before you as part of the supercomputer hearings. The subject that I'm addressing myself to, insuring the preeminence of the United States supercomputer industry, is vital to our Nation.

The scientific and economic preeminence of the U.S. supercomputer industry depends on the strength, vigor and dynamism of our unique entrepreneurial American culture and heritage. The United States has maintained and will continue to maintain its preeminence because of our abilities to encourage the emergence of the high-technology venture business, expend our venture capital markets, restructure the reform of our financial and security markets, and collaborate as required on R&D systems between our private and public sector institutions.

In my opinion, American strength has always been our ability to be scientifically creative, technologically adept, and managerially innovative. In these ways, we have met, and will continue to meet, critical challenges in order to promote the common good and general welfare.

We are now engaged in a fierce computer scientific and eocnomic competition.

To conserve time, I have a few slides.

[See slide 1, page 16.]

Dr. **Kozmetsky**. The first slide—why don't you focus it, J.R.—simply tells you that I shall be covering very quickly some historical developments, come immediately to the present, forecast as best as I can the future, and then discuss with the remaining time the economic consequences of losing pre-eminence in the supercomputers.

As you well know, historically, the computer has been viewed as one of the most basic tools of the last third of the 20th century. Computer usage for defense, space exploration, and industry, education, manufacturing, and research and in everyday life, underpins the opportunities of our information knowledge age.

[See slide 2, page 17.]

Dr. **Kozmetsky**. The next slide, the table shows the historical leadtimes between the first working models and the first commercialization by nations in the digital computer industry.

A careful look at that—you will notice I have the United States on top in the green. It took us about 6 years to commercialize our first one and, as of 1983, our technology is dominant and our market position is worldwide.

If you look at the bottom, you will see that Japan was able to commercialize its first computers in 2 years. Currently, it's major in technology development, and it certainly has a worldwide market.

This table reflects two types of government policies toward the indigenous computer industry. One is to encourage the use of computers without R&D and production capabilities, and the other is to promote the development of full-scale computer industry focused on both the domestic and international markets.

Up to the 1970's, it was perceived that the proper use of computers could accelerate the development of a nation's entire economy, rather than concentrating on the R&D and production capabilities to promote one industry, namely the computer industry. This was true in the case of Austria, Switzerland, and Australia, which are in orange. And you will see that they're not dominant in technology nor in market position.

Other major European nations, such as France, the United Kingdom, West Germany, Italy, Sweden, Denmark, and Belgium—as well as Japan—adopted a policy to foster a computer industry. They believed that they had ample resources for developing the entire scientific and engineering infrastructure, as well as electronic components, computer equipment, peripherals, and required software. This policy saw the computer industry as part of the nation's industrial complex that would become an important force in their nation's development.

By the end of the 1970's only three nations had any significant computer industry and market, the United States, Japan and the United Kingdom. The United Kingdom had 60 percent of its own domestic market.

[See slide 3, page 18.]

Dr. **Kozmetsky**. We'll go to the—the critical thing throughout all of this is what's been happening in the 1980's, so let me—and they've sort of been unrelated in the literature and in discussions, which I think are dramatically altering the U.S. computer industry.

The first of them has been the maturation of our own companies. By the end—before the end of the seventies, we saw a shakeout in our first-tier computer companies; that is, we saw GE and RCA leave the market, and IBM became dominant. I know there are a lot of friends here in the second-tier computer manufactures, so I hope I don't mind me calling them second tier.

But they remained and selected specialized data processing market segments. And certainly in the scientific and engineering supercomputer, the Government market dominated and it was completely held by CDC and Cray.

But of more importance, what's been happening in the last few years since 1976, 1977, is to really realize what the competitive impacts of the IBM system philosophy on large-scale computers has been. First, IBM deliberately started out to make its software compatible for its own future generations, which then said that billions of dollars that have been invested in software could be carried forward and not redone.

Second, this permitted, along with their standardization of their CPU's and peripherals, were the interfaces. Both the peripherals and central processing could also be areas subject to competition from other companies.

The next was that the IBM compatible market emerged. And with that we had the following consequences. That the computer competition now shifts to the way you design the computer and advances that you make in science and advanced engineering.

[See slide 4, page 19.]

Dr. **Kozmetsky**. Second, we saw a decided drop in prices of computers that we had not seen before.

The other most amazing thing was to see a completely larger market than most of us in the computer industry ever envisaged, almost an order of magnitude. The elasticity of demand for computers was larger.

At the same time, there has been ongoing, as the chairman made in his opening remarks, a whole set of large-scale demands for supercomputers that we had not thought of before. Amongst them were the Government demands, of which he mentioned three. But they're also undergoing, in various areas of the Federal Government's national policy modeling, including various committees of the Congress, and also just simply getting more productive and efficient in data processing, like the IRS, social security, and other large scale.

There has been an explosion of research for scientific computation. But the most important thing that emerged is that we can tell you unequivocally in our research at the university that no longer is either a large supercomputer by itself sufficient, but we do need to have also a large-scale data-processing computer. We need both.

Now, these emerging needs also—if you will go to the next slide—comes at the same time that we see two major technologies coming together, computer technology and the communications technology.

[See slide 5, page 20.]

Dr. **Kozmetsky**. Now, the impact of that you can study a little later, but what it's doing is suddenly giving birth to a 1993 industry, which I have estimated to be between $500 and $700 billion a year, where they are currently, in my estimation, ranging somewhere from $100 to $150 billion a year. So the coming fierce competition is extremely important, especially when you view that the United States is still 40 percent of the total world market.

Now, what we're really in and what we've been seeing happening in our country to meet the so-called supercomputer competition from Japan is on the next slide.

[See slide 6, page 21.]

Dr. **Kozmetsky**. And I have decided to say what we really are talking about is not so much economic competition, as global scientific competition, that really most of the programs that we see in the United States are really approaches for scientific preeminence rather than economic.

If we look closely at MCC, it's to develop new technology, new methods, which that, in turn, are licensed to those who are cooperating and part of that organization. If we look at Stanford Semiconductor Research Corp., we see that, too, is scientific and technological research and development, and that, too, is to be picked up by the industry somehow and to be used. Microelectronics Center of North Carolina also is very actively involved in advancing the sciences.

And, finally, DARPA itself is advancing the sciences, and its commercialization is left to traditional methods of spinoff that somehow the current laws and ways we've been doing things, that it will get into the commercial market.

This is quite in contrast to the Japanese strategy. The supercomputer program is their calculations, and they have decided to target that where they give grants of 100 percent to industry. Currently 30 Government researchers have been sent to Fujitsu's laboratories. And the other companies in the industry have sent a minimum of two engineers. They are working there. And they have a target, the date, and the supercomputer targets are very much like the American out in the 1990 timeframe.

When it's finished, Fujitsu will have the laboratory, the Government people go back, and know-how will be easier to transmit than simply trying to go around the wide group of geographically dispersed research.

The fifth generation, many of us will discuss will it ever get into being as a machine. The real critical thing for this hearing is that they, too, are realistic, and it said that they're going to program spinoffs from the machines and research that they're doing that they think will pay for it. The United Kingdom has just recently decided to put in $550 million to also go into the supercomputers. Currently, they're matching one-half of whatever the industry puts up, and 100 percent of what the university puts up and, on a total, they'll provide about 80 percent of the $550 million.

Europeans are somewhat behind. They are discussing putting in ESPRIT, and they are still discussing it and want to see it for 5 years out. And then things are up in the air yet.

In the meantime, Bull of France, ICL in England, and Siemens have established a joint research institute located in Bavaria, and they're taking off to do joint industry research.

[See slide 7, page 22.]

Dr. **Kozmetsky**. The next chart is kind of a way of seeing how we can pool all this together.

If you look at the very bottom even line, that's the Japanese funding of supercomputers, and you'll see it's a nice steady line. It does not include any research and development funds from their companies. I've estimated that if we took the standard amount that their companies have been putting into R&D, they will be larger by about $40 million a year.

For the top line, which is a combination of everything that the United States has going, both in its private companies and the various Government assistances that are going on and its various States that are doing. The European has a sharp dropoff after 5 years. And, as you can see, England is at a lower level and goes off. So it seems to me as we end the eighties, we're going to be back in a competition between two.

[See slide 8, page 23.]

Dr. **Kozmetsky**. Just to show you some numbers so that you get some idea of the total U.S. R&D, it's someplace at least $3.3 billion over the next decade. On the other hand, the most weakest number that I have up there is the top one, which is the supercomputer corporate R&D. I've just simply taken the highest numbers they have ever used, I have projected their sales for the next decade, and I think that that's way too high. But even if you cut it in half, we're still putting close to $3 billion in research.

All right.

[See slide 9, page 24.]

Dr. **Kozmetsky**. Now, let's look more at the present. The computer industry's worldwide market is no longer as it simply was. Today it is a complex competitive environment which both the United States and Japan exert major influence across all segments. Both of—the emerging computer and communications industry will include technologies for so many emerging industries, as it's being talked about. So that it's extremely important that we realize that the supercomputer is the heart of so much of the emerging industries that people today refer to as the coming high-tech industry.

I would like to share with you some of the data which I have. Fujitsu this month has delivered its first supercomputer to its university. It has also has done something that I'm not aware that any other manufacturer in the United States has done, and that is to make it compatible with the IBM large-scale computer.

I already know that—it has not been publicized—but some of the non-Government applications in the United States have already been transmitted from the United States through satellite into Tokyo, run on their vector computer, and it has come back with—in less than 24 hours, and has been benchmarked as accurate as the ones that—of—their data that they're getting off of the American computer. And this adds a new dimension.

So I think that, at the present moment, as I look at it, it says our scientific—that's the slide 5. You've got it? No, you went to the charts. That's all right. Hold it back. I am not going to forecast yet. Jake, go back, J.R. Just hold it up. Hold that one up. All right. Turn the light off. Good.

[See slide 10, page 25.]

Dr. KOZMETSKY. Let me just summarize where I think we are.

The scientific supercomputer has become a worldwide regional competition for emerging global markets. Two of the scientific competitors do not have major Government funding, the United States and the European nations. Defense R&D, NSF, DOE, support our own competitors.

Japan has extensive support in terms of funds and research scientists. They are also planning spinoffs. They've also targeted markets, while the U.S. firms target incremental technological breakthroughs. Only Japan has firm, dedicated funding for 10 years. The U.S. firm funding is dependent upon the market and the individual firm's ability to generate earnings.

Now if, as I look to the future, the supercomputer industry in this country follows the same policies as it did in the sixties and seventies, that is if it builds a flourishing similar computer industry rather than a focus on the use of the supercomputer and its technologies, then I can forecast the results.

Forecasting is an integral part of the innovation process. It is action and policy oriented. Its major assumption is that one can direct the path into the future which can be either self-fulfilling or self-defeating. Therefore, it is necessary to forecast supercomputer annual sales, their prices on the high, medium and low side, and then project the annual world supercomputer sales.

Because of the compatibility that the Japanese now have with the large-scale IBM computer, I would like to show you my projection of worldwide computer sales.

[See slide 11, page 26.]

Dr. KOZMETSKY. Most of the American literature and guesstimates and forecasts have been that we could not sell more than 50 a year, or maybe a total of 500 through this decade. Because of the compatibility nature now that supercomputers can have with the other large data processing, I have estimated that, in this decade, we'll sell 5,000.

[See slide 12, page 27.]

The next thing was how could I forecast the price for most people in the computer industry—I can assure you the Japanese do use learning curve prices. I use three different learning curve rates, at the 80, 90 and 85 percent, 90 being the low—or the high prices, if you want—80 the low and 85 in the middle.

[See slide 13, page 28.]

Dr. KOZMETSKY. That then says that, if you look at the next chart, that by the end of 1993, there could be anywhere from $2.2 billion to $6.4 billion of sales.

Now most of you will think that $2.2 billion or $6.4 billion isn't large. So I thought the easiest way to relate that is to show you my projections of IBM revenue and R&D expenditures. You show that one.

[See slide 14, page 29.]

Dr. KOZMETSKY. You'll see that IBM is currently slightly above 40 for 1984, and it will hit someplace above $170 billion by the end of 1993. Its current R&D is ranging close to $3 billion, and it will be over $10 billion by 1993.

So your first feeling is either $2.2 or $6.4 billion of supercomputers is small compared to $170 billion. So I'd like to put that in relationship to see what its meaning is.

It will provide, if we gut the whole market in this country, between 80,000 and 107,000 direct jobs, with another 100,000 to 250,000 people to operate these supercomputers. Direct employment is equal to the total number of people we have employed in furniture and fixtures in the United States, as well as what the Department of Commerce calls miscellaneous marketing. To put it in the Texas perspective, that's one and a half times more employment than in the petroleum and coal production industry, also one and a half times more than we have employees in the leather and leather production. So those are very significant numbers.

Now let's see how fast the Japanese can move. I've said they're already in production. Let's sort of take a look at what they did to us on the 64K RAM.

[See slide 15, page 30.]

Dr. KOZMETSKY. And as you all know, they took away 70 percent of the market between 1979 and 1982 from us. My forecast of the semiconductor market for 1983 was $7 billion in the United States, on its way to $45 billion, or the total for the world is $16 billion to $90 billion.

The Japanese think that they can maintain a minimum of 50 percent of the market. Well, let's take a look at the next slide.

[See slide 16, page 31.]

Dr. KOZMETSKY. If they do, American chip manufacturers will move their production from $3.5 billion to $11.25 billion, while the Japanese move theirs from $1.2 billion to $11.25 billion. And those are very dramatic results. In the short, the Japanese will get 10 times more of the market by 1993, while we're getting about three times for domestic.

If you look at that chart, you'll also see that U.S. internal company operations will increase dramatically. And in my formal paper I said I see signs of the same problems that we have in the steel industry, dumping.

[See slide 17, page 32.]

Dr. KOZMETSKY. Now if I went to the next chart, let's see if I can't break down the forecast of supercomputer sales in terms of which two of the dominant nations might be able to get.

The first chart there shows annual unit sales. And on the bottom line is noncompatible supercomputers with IBM, and I have assumed the American manufacturers take all that. And that IBM compatible market goes to Japan, and that is the top curve.

[See slide 18, page 33.]

Dr. **Kozmetsky**. Translating that into dollars, it says by 1993, the Japanese could surpass our domestic companies' sales by 1988 and, by 1993, end up with $2.6 billion of sales versus ours at $1 billion. For the manufacturers in the room, I use 90 percent for American and 80 percent for the Japanese.

[See slide 19, page 34.]

Dr. **Kozmetsky**. Another way of showing what would happen total marketwide, taking numbers off those charts, on this table it says in 1983, the Japanese supercomputer market was just $0.01 billion, by 1993 it would be two and a half, while ours going from $500 million to $1 billion per year.

Now, in this setting, there are a number of economic consequences that can be evaluated. First, the global computer industry during the eighties can be progressively less dominated by U.S. companies. The 1990's will be even more competitive than the 1970–80 era.

This will not be the result of the loss of scientific preeminence. Rather, it's a natural progression to lose the share of global markets when other nations become directly competitive.

The key question is, what share of the future computer communication market can or should the United States realistically expect to maintain? The global market for computer communications is estimated to increase from $100 to $150 billion in 1983 to $500 to $700 billion in constant 1982 dollars in 1993. Currently, the United States' scientific and technological preeminence resulted, in our estimation, to be at least 75 percent of the current market, or $75 to $120 billion.

If U.S. preeminence in its selected fields continues, its shares could drop to 50 to 60 percent. Therefore, in 1993, we would have between $250 to $400 billion in constant 1982 dollars.

Other competitive nations, including Japan, the United Kingdom, and the European combinations, could increase market share from $25 to $30 billion to $250 to $280 billion under these assumptions. Economically, all would prosper, even though scientific and technological preeminence was shared. In short, it's difficult to imagine the United States maintaining its current market share over the longer run.

Now the U.S. supercomputer industry share of worldwide markets could be dramatically below the 50 percent, as I showed in the charts. IBM or others must enter the market for the United States to maintain 50 percent or more. IBM could be well ready to announce soon its supercomputer. But there is no overt evidence available at this time.

The loss of American market share could have a relatively immediate economic consequence in the form of price wars for computers and selected semiconductor chips during the next 5 years. This could result in less R&D funds for U.S. companies to continue developing the next generation of supercomputers.

[See slide 20, page 35.]

Dr. **Kozmetsky**. If the prices go down dramatically, and they come right off the top, you'll see that our percent of revenue—Cray's, the big one there, that spends 20 percent of its revenue on R&D—all the others are below 11 percent. So at 10 or 15 percent strategic drop in prices could dramatically reduce the amount of funds our supercomputer companies would have to maintain preeminence.

Now the supercomputer scientific competition makes clear that the volume of resources, including human talent, devoted to fundamental science and to its technological innovations is limited. In simple supply and demand terms, this means that science and engineering human resources, when scarce, will be diverted to increasing returns on investment, rather than on maintaining scientific preeminence.

The choice in pure economic theory would be to select those demands which provide the highest marginal use. According to managerial theory, this would mean selecting those technological products and services that provide the highest gross margin, as well as rate of return on investment commensurate with liquidity.

Since the results of science do not enter directly into the private economy, the demand for prestige and academic excellence may require public support under such economic conditions. To maintain scientific preeminence on a significant economic share of the supercomputer market, the U.S. Government should consider ways of supporting the domestic markets for American supercomputers. One way to support the domestic market, while increasing the number of scientists, technologists, and engineers is for the Government to assist in placing supercomputers in major universities and nonprofit research institutions across the country.

Mr. Chairman, regardless of all the fine laws you helped pass to get better education of science and engineers, we today have such a shortage, and the less time that we have between generations, that we're really hurting. As one president of a major computer company told me, he's prepared to invite foreign countries to send him their scientists and engineers while they develop the next model, then they can go back and compete, because they're just that short.

So a more rapid dispersion of supercomputers than the market currently dictates will help to insure the transfer of technology which can support widespread continued R&D activities that will expand the whole computer communications industry. At the same time, this dispersion will expand the knowledge and experience available to our human resource base. Supercomputers are still not available in most of our prestigious academic institutions.

In addition, the supercomputer is needed for other scientific developments in key areas, in addition to defense and energy and meteorological forecasting. We desperately need it in space, astronomy, biology, chemistry, genetics, biotechnology, and national policy simulation and assessment, and large-scale project management.

Now the viability of our computer communications industry depends not only on maintaining, but also expanding, private sector R&D funding through revenues and earnings. As I said, if the revenues drop 5 or 10 percent, our major computer firms will be in jeopardy and may be forced to cut their R&D activities.

The consequences of losing economic and scientific preeminence in the supercomputer industry are vast. The supercomputer is a central driver for the rapidly emerging worldwide computer communications industry. It impacts communications development, the renewal of basic industry, productivity increases, and the development and expansion of new industries. It's essential in improving our educational structure, fulfilling critical manpower structures, and enhancing our industrial creativity and innovation. It is the seed for encouraging the emergence of a myriad of technology venture businesses in the context of a private enterprise system that always has been the unique American way to achieve and maintain U.S. economic and scientific preeminence.

Those are my prepared remarks, Mr. Chairman.

[The slides, biographical sketch and prepared statement of Dr. Kozmetsky follow:]

OVERVIEW

I) HISTORICAL DEVELOPMENTS.

II) THE PRESENT.

III) THE FUTURE OF THE SUPERCOMPUTER INDUSTRY: A FORECAST.

IV) THE ECONOMIC CONSEQUENSES OF LOSING PREEMINENCE IN SUPERCOMPUTERS.

Slide 2

TABLE 1

Historical Lead-Times and Current Status
in the Computer Industry

By Selected Nations

Nation	First Working Model	First Commercialization	Current Technological Status	Market Position
United States	1946	1951-52	Dominant	Worldwide
Canada	1955	1964	Minor	Domestic
U.K.	1949	1953	Minor	Domestic
West Germany	1941	1954	Minor	Domestic
France	1952	1957	Minor	Domestic
Netherlands	1956	1957	Minor	Domestic
Italy	1957	1960	Minor	Domestic
Sweden	1950	1962	Minor	—
Denmark	1957	1962	Minor	—
Belgium	1960	1963	Minor	—
Austria	1959	—	—	—
Switzerland	1956	—	—	—
Australia	1954	—	—	—
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
Japan	1957	1959	Major	Worldwide

Source: Computers 1971-1981, Volume I, Center for Technlogical and Intradisciplinary Forecasting, Tel Aviv University, August 1971.

Slide 3

UNRELATED EVENTS ALTERING THE U.S. COMPUTER INDUSTRY

I) Maturation of U.S. Computer Firms

A) Shake out of first-tier computer companies in electronic data-processing market (G.E. and RCA); IBM dominant.

B) Second-tier computer manufacturers select and expand into specialized market segments.

C) Scientific-engineering supercomputer government market dominated by CDC and Cray.

II) Competitive impacts of IBM System Philosophy of Large-Scale Computers.

A) Software compatibility for future generations.

B) Interface standardization permits IBM-compatible peripherals and central processing companies to compete

C) "IBM -compatible" market segment emerged with the following consequences:

1) Computer competition shifts from components to architecture.

2) Cost reduction due to competition

3) Unforecasted increase in sales due to newly developed price elasticity.

III) Emerging Large Scale Demand forData-Processing and Computational-Intensive Computers

A) Government Demand

1. Defense- simulations
2. Energy- nuclear, petroleum demand
3. Atmospheric modeling
4. National policy modeling-many applications
5. Data storage and processing- IRS, Social Security et al.

B) Scientific computations for

1. Fusion and other energy research
2. Advanced physics
3. Chemistry
4. Genetics
5. Artificial intelligence
6. Astronomy
7. Biology
8. Advanced mathematics
9. Macro-engineering project management
10. Statistical modeling.

Slide 5

FIGURE 2

Examples of Convergence Between
Computer and Communication Technologies

1940-1990

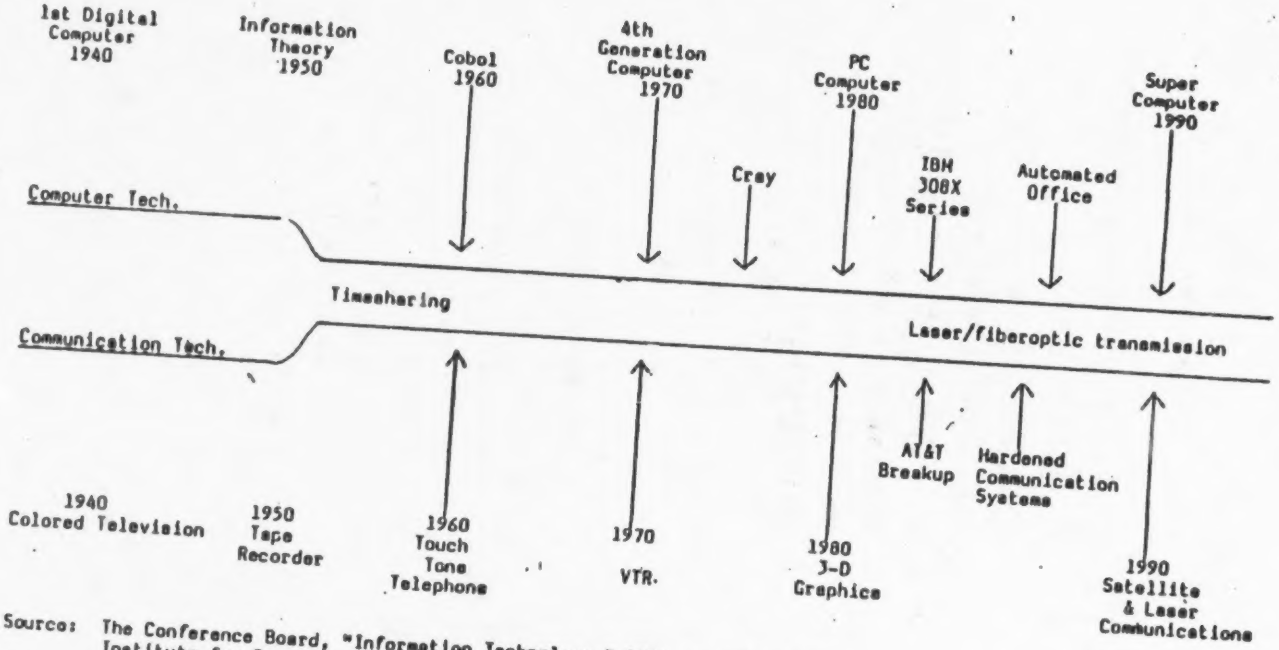

1st Digital Computer 1940

Information Theory 1950

Cobol 1960

4th Generation Computer 1970

Cray

PC Computer 1980

IBM 308X Series

Automated Office

Super Computer 1990

Computer Tech.

Timesharing

Laser/fiberoptic transmission

Communication Tech.

1940 Colored Television

1950 Tape Recorder

1960 Touch Tone Telephone

1970 VTR.

1980 3-D Graphics

AT&T Breakup

Hardened Communication Systems

1990 Satellite & Laser Communications

Source: The Conference Board, "Information Technology," 1972, p. 192 and the Institute for Constructive Capitalism, The University of Texas at Austin.

GLOBAL SCIENTIFIC COMPETITION: EMERGING APPROACHES FOR PREEMINENCE

A) United States

1. Microelectronics and Computer Technology Corporation (MCC)
2. Semiconductor Research Corporation (SRC)
3. Stanford University Center for Integrated Systems (SUCIS)
4. Microelectronics Center of North Carolina (MCNC)
5. Defense Advanced Projects Research Agency (DARPA)

B) Japan

1. Supercomputer Program
2. Fifth-Generation Computer Program

C) U. K.

1. Alvey Programme for Advanced Information Technology

D) European

1. European Strategic Program for Research and Information Technologies (ESPRIT)
2. Joint Research Institute- Bavaria

CHART A*

RESEARCH FUNDS FOR SUPERCOMPUTERS

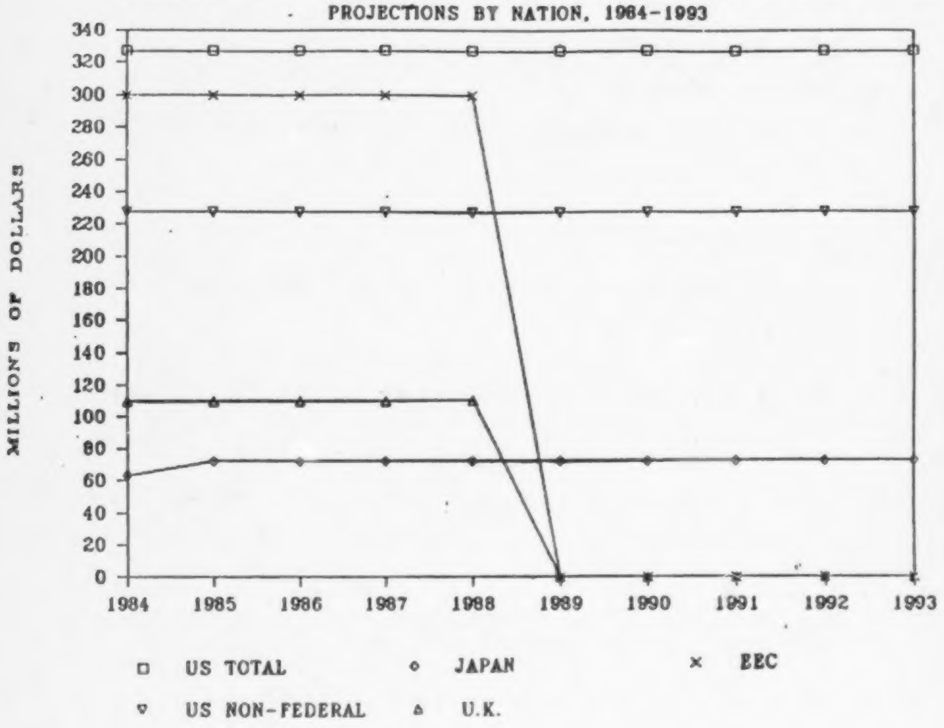

Source: Spectrum, November 1983, nad the Institute for Constructive Capitalism, November 1983.

* See following table for distribution of U. S. research funding.

Slide 8

CHART A

FORECAST
FUNDING DISTRIBUTION
U.S. SUPERCOMPUTER RESEARCH AND DEVELOPMENT
1984-1993

Research Institution	Estimated Ten-Year Expenditure in Millions of 1983 Dollars
Corporate Supercomputer R & D	$1180
Microelectronics and Computer Technology Corpotation (MCC)	1000
Semiconductor Research Corporation (SRC)	100
Stanford University Center for Integrated Systems (SUCIS)	300
Microeletronics Center of North Carolina (MCNC)	100
Subtotal-	$2300
Defense Advanced Projects Research Agency (DARPA)	1000
Total: U. S. Supercomputer Research	$3300

CRITICAL ISSUES FOR GLOBAL SUPERCOMPUTER COMPETITION

1. Scientific achievements contribute to a nation's pride and heritage.

2. Commercialization (or how to convert advances in supercomputer science & technology) is the most critical economic problem.

3. Preeminence in the market place requires the ability to manufacture at lowest cost and meet very tight delivery dates.

4. The process of designing succeeding generations of supercomputer systems is facing critical economic problems: ever shorter periods between new generations, current shortage of experienced personnel, and large capital investments that must be recouped in very short times have changed the risk-reward nature of the communications/ computer industry.

FIGURE 1 Slide 10

Japanese Policies for Promotion of Computerization

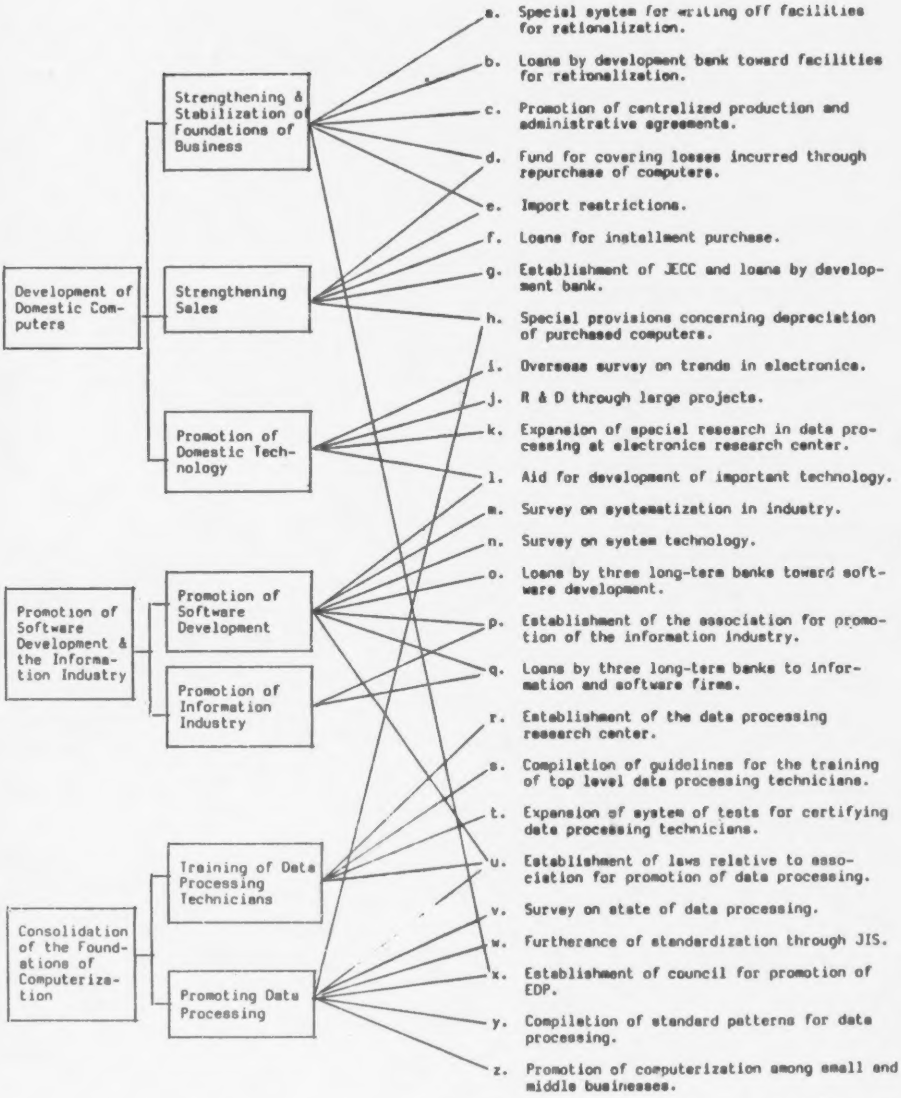
Development of Domestic Computers
Strengthening & Stabilization of Foundations of Business
Strengthening Sales
Promotion of Domestic Technology
Promotion of Software Development & the Information Industry
Promotion of Software Development
Promotion of Information Industry
Consolidation of the Foundations of Computerization
Training of Data Processing Technicians
Promoting Data Processing

a. Special system for writing off facilities for rationalization.
b. Loans by development bank toward facilities for rationalization.
c. Promotion of centralized production and administrative agreements.
d. Fund for covering losses incurred through repurchase of computers.
e. Import restrictions.
f. Loans for installment purchase.
g. Establishment of JECC and loans by development bank.
h. Special provisions concerning depreciation of purchased computers.
i. Overseas survey on trends in electronics.
j. R & D through large projects.
k. Expansion of special research in data processing at electronics research center.
l. Aid for development of important technology.
m. Survey on systematization in industry.
n. Survey on system technology.
o. Loans by three long-term banks toward software development.
p. Establishment of the association for promotion of the information industry.
q. Loans by three long-term banks to information and software firms.
r. Establishment of the data processing research center.
s. Compilation of guidelines for the training of top level data processing technicians.
t. Expansion of system of tests for certifying data processing technicians.
u. Establishment of laws relative to association for promotion of data processing.
v. Survey on state of data processing.
w. Furtherance of standardization through JIS.
x. Establishment of council for promotion of EDP.
y. Compilation of standard patterns for data processing.
z. Promotion of computerization among small and middle businesses.

Source: Japan Computer Usage Development Institute, Computer White paper, 1970 edition.

Slide 11

CHART B

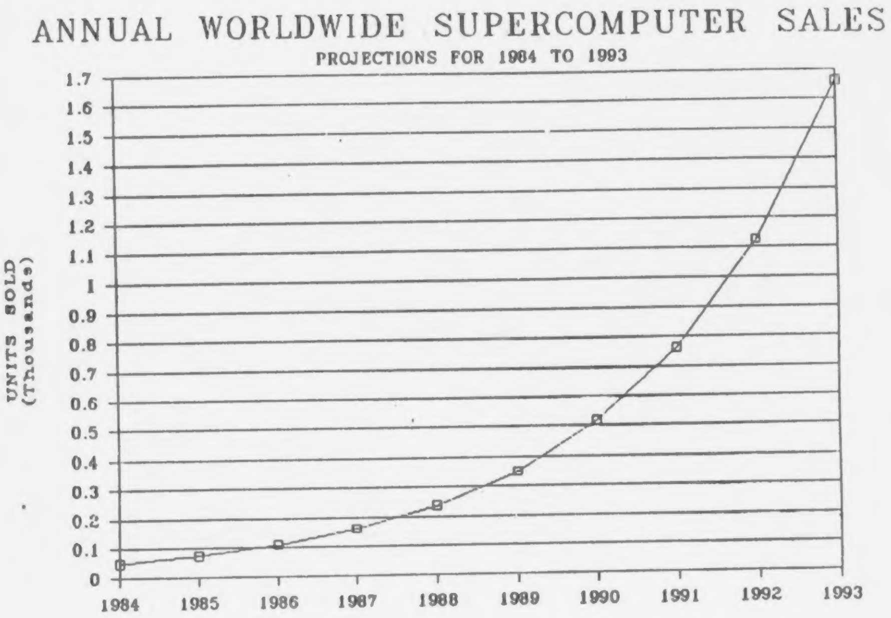

Source: The University of Texas at Austin, The Institute for Constructive Capitalism, November, 1983.

CHART C

Slide 12

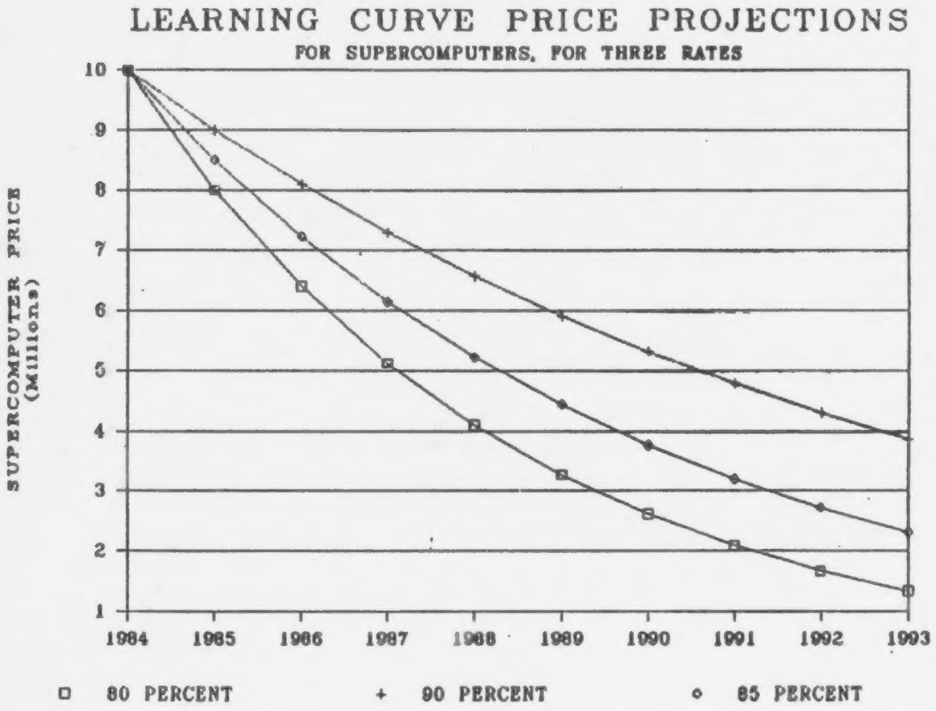

LEARNING CURVE PRICE PROJECTIONS
FOR SUPERCOMPUTERS, FOR THREE RATES

Source: The University of Texas at Austin, The Institute for Constructive Capitalism, November 1983.

CHART D

Slide 13

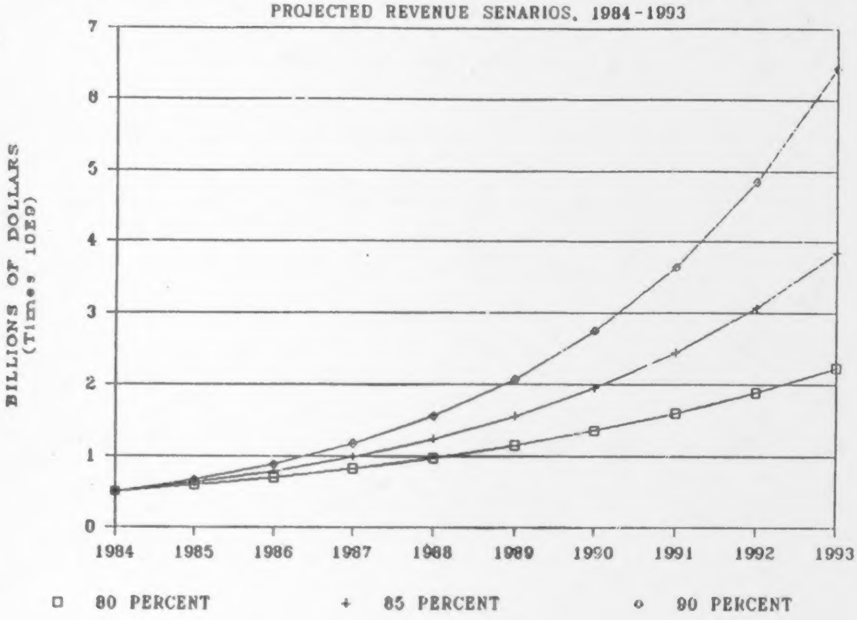

ANNUAL WORLDWIDE SUPERCOMPUTER SALES

PROJECTED REVENUE SENARIOS, 1984-1993

Source: The University of Texas at Austin, The Institute for Constructive Capitalism, November 1983.

CHART E

Slide 14

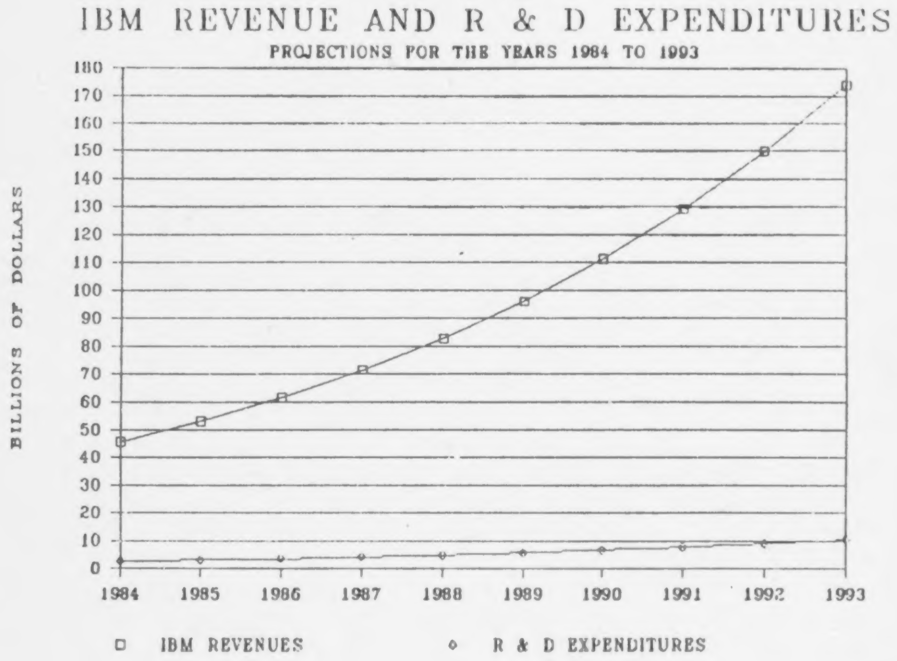

Source: First Boston, Insights, October 4, 1983, and The University of Texas at Austin, The Institute for Constructive Capitalism, November 1983.

Slide 15

TABLE 2

MARKET PROJECTIONS FOR SEMICONDUCTOR CHIPS
FOR 1983 and 1993

(In billions of 1983 dollars)

	U. S. Market Demand	Other Nations Market Demand	Estimated World Market Demand
1983	$ 7	$ 9	$16
1993	$45	$45	$90

Slide 16

TABLE 3

UNITED STATES SEMICONDUCTOR DEMAND
FOR 1983 and 1993
(In billions of 1983 dollars)

	1983	1993
Total U. S. Demand	$ 7	$45
Source of supply for Chips:		
U. S. Chip Manufacturers	$ 3.5	$11.25
Japanese Chip Manufacturers	$ 1.2	$11.25
U. S. Internal Company Production	$ 2.3	$22.5

CHART F

Slide 17

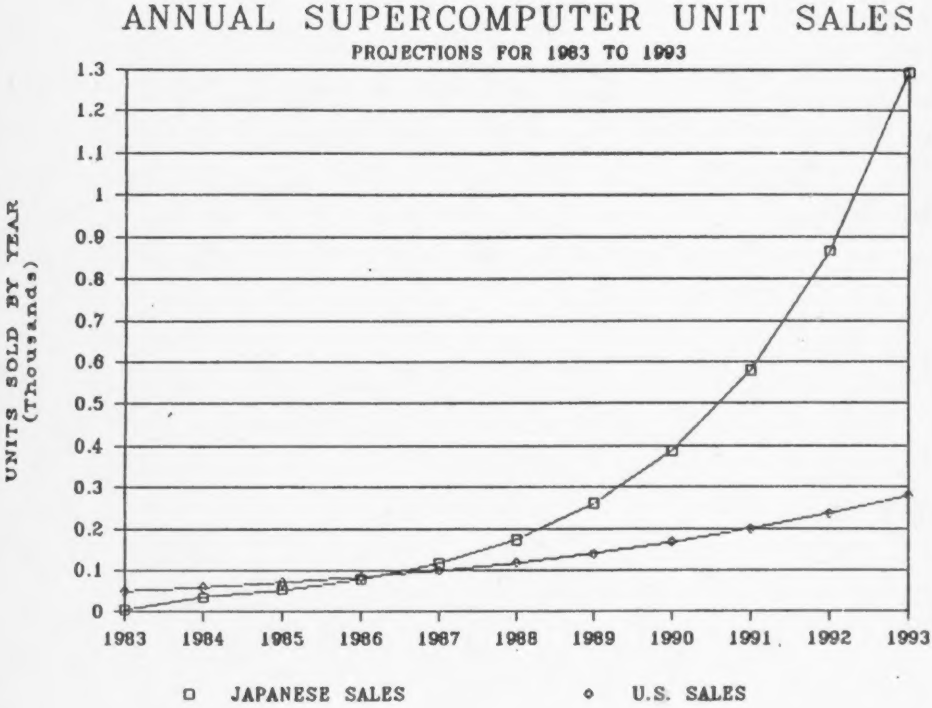

ANNUAL SUPERCOMPUTER UNIT SALES

PROJECTIONS FOR 1983 TO 1993

Source: The University of Texas at Austin, The Institute for Constructive Capitalism, November 1983.

CHART G

Slide 18

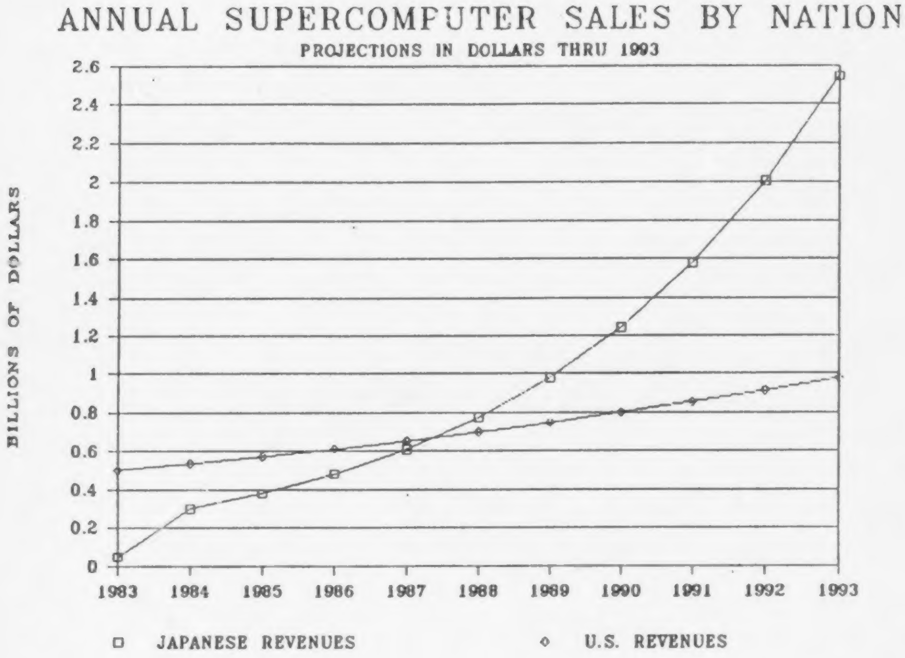

Source: The University of Texas at Austin, The Institute for Constructive Capitalism, November 1983.

FORECAST
SUPERCOMPUTER SALES
BY NATION 1983 and 1993
(in billions of 1983 dollars)

Annual Sales	Japanese Companies	U.S. Firms
1983	$ .01	$ .5
1993	$2.50	$1.0

CHART H

Slide 20

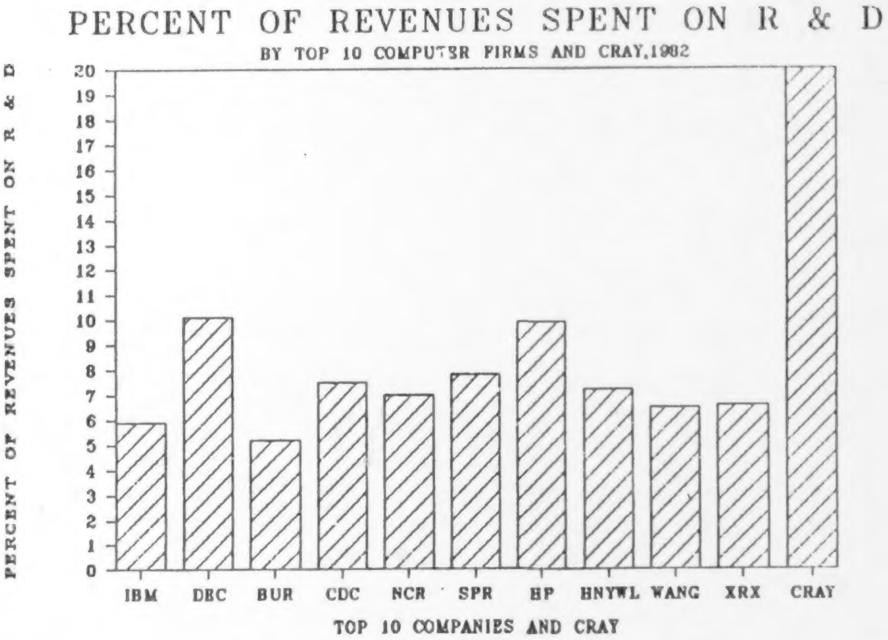

Source: Datamation, June 1983.

DR. GEORGE KOZMETSKY

Dr. George Kozmetsky is Director of the Institute for Constructive Capitalism at the University of Texas and Executive Associate for the University of Texas Board of Regents. For 16 years he served as Dean of the Graduate School of Business and Professor of Management and Computer Sciences.

Dr. Kozmetsky has a distinguished record in American business. He served in executive positions with Hughes Aircraft and Litton Industries and was a founder of Teledyne, Inc. He is currently a member of many company boards, including Amdahl and Data Point.

ENSURING U.S. PREEMINENCE IN THE SUPERCOMPUTER INDUSTRY

by

George Kozmetsky*

Testimony before the U.S. House of Representatives,
Committee on Science and Technology
November 15, 1983

*Director, Institute for Constructive Capitalism, J. Marion West Chaired Professor in Contructive Capitalism; Professor of Management and Computer Sciences, The University of Texas at Austin.

The Institute for Constructive Capitalism
The University of Texas at Austin

Introduction

Mr. Chairman, members of the Committee, I am honored to have this opportunity to appear before you as part of these Supercomputer Hearings. The subject, "Ensuring Preeminence in the U.S. Supercomputer Industry" is vital to our nation. As someone who has been intimately involved with the computer and communications industries for over three decades and as someone who is intensely concerned about the management of technology and the competitive position of the United States in an emerging, highly-competitive protection-seeking global economy, I have looked forward to testifying before this distinguished Committee.

I bring to these proceedings the personal perspective of one who has experienced through the private enterprise system the dynamic growth of the computer and communications industries and who has also sought to translate this experience into more effective academic and scientific theory.

The scientific and economic preeminence of the U.S. supercomputer industry depends on the strength, vigor and dynamism of our unique, entrepreneurial American culture and heritage. The U.S. has maintained, and will continue to maintain, its preeminence because of our abilities to encourage the emergence of the high-technology venture business, expand our venture capital markets, restructure the reform of our financial and security markets and collaborate, as required, on R & D systems between our public and private sector institutions. America's strength has always been our ability to be scientifically creative, technologically adept and managerially innovative. In these ways, we have met and will continue to meet critical challenges in order to promote the common good and provide for the general welfare.

We are now engaged in a fierce supercomputer scientific and economic competition. My remarks are directed at the consequences of losing this competition.

Acknowledgments

There are a great many people who have aided me in my understanding of supercomputers, in fact, there are far too many to list here. Nevertheless, there are some friends and colleagues whose input has been especially insightful and helpful.

From the University of Texas at Austin I would like to express my gratitude to Gerhard J. Fonken, Vice President for Academic Affairs and Research; Earnest F. Gloyna, Dean of the College of Engineering; and William Cunningham, Dean of the Graduate School of Business; to my colleagues on the University of Texas at Austin Computer Committee I thank them all, especially Professors James C. Browne, J. K. Aggarwal, Alfred G. Dale, Joe Lagowski and Charles H. Warlick. In the College and Graduate School of Business at the University of Texas at Austin, I would like to thank Dr. Eugene Konecci, The Kleberg-King Ranch Professor, Dr. Timothy Ruefli, The Fayez Sarofim Centennial Professor; and Professor Abe Charnes, The John Harkins Centennial Professor and Director of the Center for Cybernetic Studies.

In the research and administrative Staff at the Institute for Constructive Capitalism at The University of Texas at Austin, I would

like to note my especial appreciation to Dr. W. W. Cooper, Dr. Sten Thore, Dr. Robert Peterson, Dr. David Kendrick, Dr. Leon Lasdon and the Institute's Assistant Director, Dr. Raymond Smilor. I wish to thank Ms. Marilyn Kistner for her invaluable assistance in preparing this manuscript.

The University of Texas is not the only academic institution whose scholars contributed to my research. My appreciation goes out to Professor C. V. Ramamoorthy of the University of California, Berkley; Professor Edward Glaser, Former Director Jennings Computation Center, Case-Western Institute; Professor Robert M. Fano, Department of Electrical Engineering, Massachusetts Institute of Technology and Frank Davidson, Systems Dynamics Group, Massachusetts Institute of Technology.

Nor is the academic world the only place I have gained knowledge of supercomputers. I have gained many valuable insights through my association with Dr. Henry Singleton, Chairman of the Board and CEO, Teledyne, Inc; Ronald V. Hill, former President of OMEX; Gordon E. Moore, Chairman of the Board and CEO of Intel Corporation; and Michael D. Gill, Jr. Assistant to the Chairman of the Board, MCO Resources, Inc.

I. Historical Developments--A Background for Understanding Supercomputer Developments

The computer has been viewed as one of the most basic tools of the last third of the 20th century. Computer usage for defense, for space exploration, in industry, in education, in manufacturing and in everyday life underpins the economic opportunities of our information/knowledge age. Through the 1970's, there were three basic segments that prescribed the computer industry:

1. R & D and production of electronic components;
2. R & D and production of hardware--computer mainframe and peripheral equipments; and
3. development and production of operational software.[1]

Table 1 shows the historical lead time between the first working models and first commercialization by nation in the digital computer industry. Their technological status as of 1983 and market position are also indicated in Table 1. This table does not take into account the military markets for digital computers. Status as to technology is indicated as dominant, major and minor. Market position is indicated as worldwide, domestic or none.

Table 1 reflects two types of government policy toward the indigenous computer industry. One is to encourge the use of computers without R & D and production capability. The other is to promote the development of a full scale computer industry focused on both domestic and international markets.

---

[1] Kahn, H. and A. Weiner, "The Year 2000," p. 21.

TABLE 1

Historical Lead-Times and Current Status in the Computer Industry

By Selected Nations

Nation	First Working Model	First Commercialization	Current Technological Status	Market Position
United States	1946	1951-52	Dominant	Worldwide
Canada	1955	1964	Minor	Domestic
U.K.	1949	1953	Minor	Domestic
West Germany	1941	1954	Minor	Domestic
France	1952	1957	Minor	Domestic
Netherlands	1956	1957	Minor	Domestic
Italy	1957	1960	Minor	Domestic
Sweden	1950	1962	Minor	--
Denmark	1957	1962	Minor	--
Belgium	1960	1963	Minor	--
Austria	1959	--	--	--
Switzerland	1956	--	--	--
Australia	1954	--	--	--
U.S.S.R.	1953	1958	Minor	East. Bloc
East Germany	1955	1964	Minor	East. Bloc
Poland	1959	1965	Minor	East. Bloc
China	1958	1966	Minor	Domestic
Japan	1957	1959	Major	Worldwide

Source: Computers 1971-1981, Volume I, Center for Technological and Intradisciplinary Forecasting, Tel Aviv University, August 1971.

Up to the 1970's, it was perceived that the proper use of computers could accelerate the development of a nation's entire economy rather than concentrate only on the R & D and production capabilities to promote one industry, namely, the computer industry. This was particularly the case in Austria, Switzerland and Australia.

Major European nations such as France, the U.K., West Germany, Italy, Sweden, Denmark and Belgium, as well as Japan, adopted a policy to foster a computer industry. They believed that they had ample resources for developing an entire scientific and engineering infrastructure as well as electronic components, computer equipment, peripherals and required software. This policy saw the computer industry as part of a national industrial complex that would become an important force in their nations' economic development.

By the end of the 1970's, only two nations besides the U.S. had developed significant computer industries: Japan and the U.K. Japan's first commercialization occurred six years after that in the U.K., and 7-8 years after the U.S. While Japan was dependent on U.S. imports through the 1970's, it has achieved independence beginning in the 1980's. An outline of Japanese policies for promotion of computerization is shown in Figure 1. These policies recognized that the transformation to a computer industry required an investment that was not profitable in the short-term, e.g., 1-5 years. They also acknowledged the need to invest in foreign technologies (i.e., U.S.) to develop an advanced electronic components industry and semi-automated production lines for manufacturing and to promote R & D for architec-

FIGURE 1

Japanese Policies for Promotion of Computerization

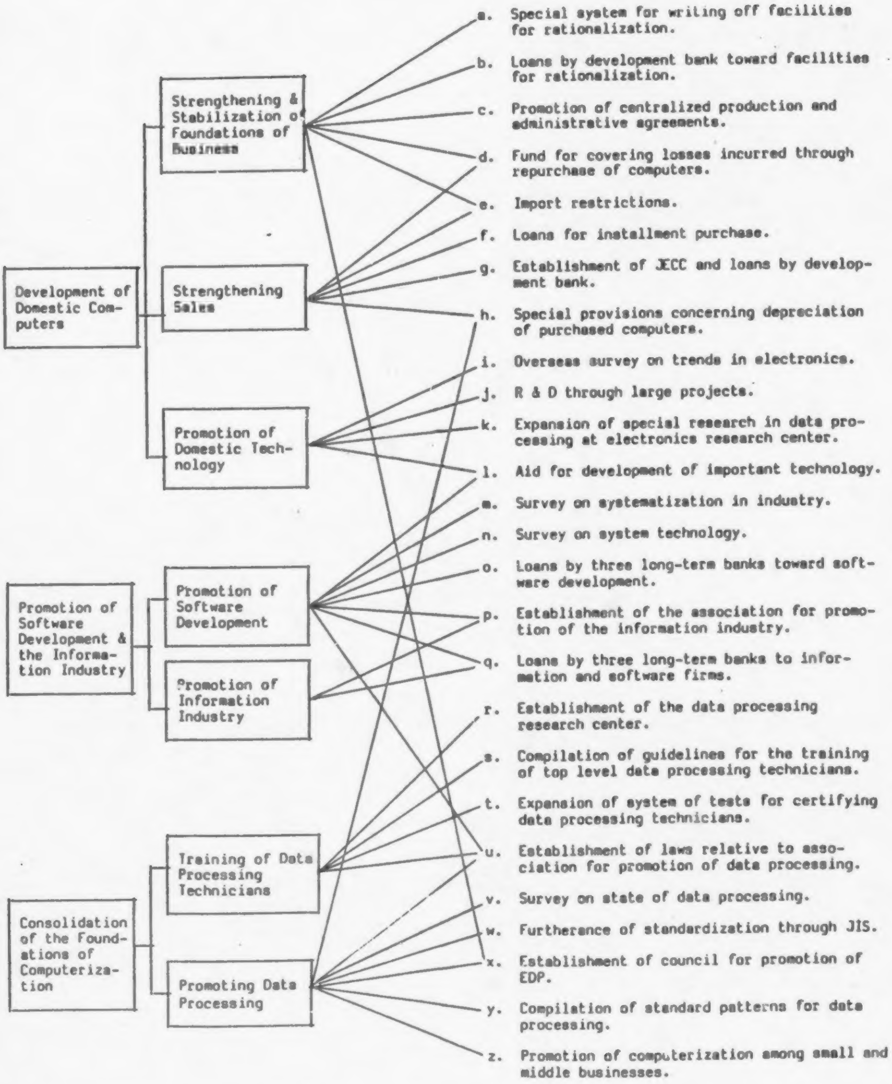
Development of Domestic Computers
Strengthening & Stabilization of Foundations of Business
Strengthening Sales
Promotion of Domestic Technology
Promotion of Software Development & the Information Industry
Promotion of Software Development
Promotion of Information Industry
Consolidation of the Foundations of Computerization
Training of Data Processing Technicians
Promoting Data Processing

a. Special system for writing off facilities for rationalization.
b. Loans by development bank toward facilities for rationalization.
c. Promotion of centralized production and administrative agreements.
d. Fund for covering losses incurred through repurchase of computers.
e. Import restrictions.
f. Loans for installment purchase.
g. Establishment of JECC and loans by development bank.
h. Special provisions concerning depreciation of purchased computers.
i. Overseas survey on trends in electronics.
j. R & D through large projects.
k. Expansion of special research in data processing at electronics research center.
l. Aid for development of important technology.
m. Survey on systematization in industry.
n. Survey on system technology.
o. Loans by three long-term banks toward software development.
p. Establishment of the association for promotion of the information industry.
q. Loans by three long-term banks to information and software firms.
r. Establishment of the data processing research center.
s. Compilation of guidelines for the training of top level data processing technicians.
t. Expansion of system of tests for certifying data processing technicians.
u. Establishment of laws relative to association for promotion of data processing.
v. Survey on state of data processing.
w. Furtherance of standardization through JIS.
x. Establishment of council for promotion of EDP.
y. Compilation of standard patterns for data processing.
z. Promotion of computerization among small and middle businesses.

Source: Japan Computer Usage Development Institute, Computer White paper, 1970 edition.

ture of computer systems. In summary, Japan invested in the requirements for a computer industry contrasted to the use of computers for the development of an economic infrastructure. By 1970, there were six Japanese companies in the computer industry.

England, on the other hand, adopted national policies to achieve a rapid increase in computers and computer techniques in industry and commerce plus "a flourishing British computer industry."[2] By the end of the 1970's, the U.K. was providing 60 percent of their own computer needs for the British government and manufacturing markets for electronic data processing. Their industry consisted of seven computer companies.

The dominant computer industry by the end of the 1970's was in the United States. There were at least a dozen major computer companies, hundreds of small and medium computer firms, a half dozen dominant electronic component companies as well as hundreds of small and medium component firms, hundreds of software firms, over a thousand service bureaus, over 75 computer rental and lease firms and several thousand by-the-hour computer rental companies. In addition, there were hundreds of universities and colleges teaching circuit designs as well as computer sciences. In other words, by 1980 the U.S. could boast a diverse and entrepreneurial computer industry that did not require large scale, direct government support. The industry responded to domestic and internatinoal market opportunities. It also responded to productivity through lower labor costs by manufacturing overseas, particularly for electronic and other components.

---

[2] Benn, A. W., Data Systems, March 1969.

During the 1970's, there were a series of seemingly unrelated events that served to alter the U.S. computer industry and thus change its basic character. By 1983, their impacts transformed the worldwide computer/communications markets and created new needs for review of U.S. governmental policies. The series of events are enumerated below:

1. The U.S. domestic electronic data processing market matured. It became a market for other nations to sell into.
2. The first tier large scale general purpose digital computer manufacturers were shaken out, namely G.E. and RCA. IBM became the dominant research and development manufacturer and marketer.
3. The second tier, large scale manufacturers, (e.g., Digital Equipment Corp, Control Data Corp., Burroughs, NCR, Sperry and Honeywell) selected and then expanded into specialized segments of the data processing markets.
4. CDC became the dominant U.S. scientific-engineering computer company in manufacturing and marketing. CDC has since spun-off its supercomputer division into a separate entraprepreneurially run company named Ada Systems. In the 1970's, two new firms emerged (Cray & Denelcor) for strictly the scientific supercomputer market.
5. The supercomputer market was created by needs fostered by

Federal Government agencies for nuclear defense and energy research.

6. In 1964, IBM developed for the first time a large scale computer system philosophy (i.e., System 360) that provided for compatibility of software between its own future generations of computers. Namely, software programs developed for one model could be run on other, more sophisticated, future models without modification. The enormous competitive impact of this compatibility was not generally realized until the late 1970's. Thus, IBM may have established data processing industry standards through computer operational compatible software. This made it possible for other CPU manufacturers to compete.
7. IBM strategy with System 360 also established a way of standardizing computer hardware, both with peripherals and the CPU. This saw the start-up of a number of competitors in both the U.S. and Japan.
8. The IBM System 360 was the catalyst for a whole chain of events in the electronic components field. Because of peripheral interface and software standardization, it became possible for competitors to use components different from that of IBM. This changed the technical and economic performance of the computer market. In turn, it transformed the economic competitive structure of the computer industry in the U.S. and the world.

9. The IBM System 360, from a large scale computer point of view, extended the competition to machine design, e.g., computer architecture. In short, it added scientific competition to the computer industry. This permitted and encouraged creativity and innovation in the computer industry worldwide.
10. Capital investments by the user changed dramatically. The billions of dollars invested in software were not obsoleted by changes in future generations of computers.
11. Competition between IBM and other IBM compatible manufacturers resulted in dramatic reductions of costs of large scale computer systems and increased the through-put performance of IBM's 370 and 308X systems dramatically. More importantly, it disclosed that the elasticity of demand for large sale computers was greater than anyone had conceived, e.g., in the order of one magnitude. In short, there was a far bigger market in the U.S. for large scale computers than originally anticipated or forecasted.
12. Concentration on scientific/engineering computation for government applications, such as nuclear and defense purposes, limited the number of machines to those restricted government and other laboratories involved in security applications.

    OPEC-induced oil shortages gave rise to a need for energy modeling and computations. Nuclear, defense and energy modeling applications created an initial market for the supercomputer. This resulted in the first working models of supercomputers and their first commercialization before the 1980's. By 1983, the market was fundamentally domestic with the beginnings of an international market dominated by three U.S. companies.

13. In the 1970's, a series of Federal Government needs for very large scale data and information processing were addressed. Among them were national policy simulations, Presidential reporting systems, early warning systems for technology assessment, material and supply shortages and other large scale modeling efforts. At the same time, national security giant data evaluation problems arose. Cost reduction and increased productivity needs in the Department of Justice, IRS and Social Security System through computerization were recognized. At first, many of these applications seemed to be more concerned with data processing and storage than with number calculations. But by 1983, it became evident that both data processing for decision making, data storage and mathematical manipulation were intimitely intertwined. This gave rise to the need for supercomputers along the lines of the Japanese 5th Generation as well as the scientific/engineering supercomputer.

14. A number of major universities in the 1970's were involved in advanced physical, natural and social science research that required both large scale data manipulation and scientific calculations for fusion, advanced physics, astronomy, biology, chemistry, genetics, biotechnology, mathematics, statistics, artificial intelligence, atmospheric physics and meteorology, water system engineering, various engineering/business large scale modeling, methodologies for macro-engineering project management, etc.

II. The Present

The global computer industry in 1983 is still dominated by the U.S. IBM can be given much of the credit for maintaining American global computer industry preeminence. Because of the range of events that took place in the 1970's, however, two nations entered the decade of the 1980's as preeminent worldwide in the computer industry--the U.S. and Japan.

The computer industry worldwide market is no longer as simple as it once was. Today it is a complex competitive environment in which both the U.S. and Japan exert major influence across all segments of the computer industry, which includes the communications industry. Indeed, the marketplace is being transformed by the fusion of two technologies--computers and communications. Both of these technologies are coming rapidly together. This convergence is changing the nature and competitive structure of each industry. The emerging computer/communications industry will include technologies for applications in areas such as local and long-distance networking, timesharing, satellites and laser and fiberoptic transmission.

An examination of current developments in each of these industries is important. The market for the computer industry is expanding day by day. It includes mainframes, the microprocessor, the home personal computer, the professional computer and a wide range of application software. The communications market is also changing in fundamental ways. Amongst the changes are the break-up of AT & T into local and long-range networks, the development of satellites, digital switching, cable and mobile systems and the rise of telecommunications. This emerging two-technology industry can be seen as giving birth to competitive markets that were not previously available, namely:

1. the automated office;
2. distributed networks;
3. automated data bases and information acquisition services;
4. basic and applied research services;
5. graphics for industry and commerce; and
6. software for the leisure/media industries.

The application software systems as well as databases, information bases and knowledge bases that service these markets are also expanding.

Figure 2 illustrates the convergence of computer and communications technologies and the resulting fused technologies for the future. It emphasizes the fusing of the two lines of developments and shows the interrelatedness of the two industries.

To understand the economic consequences of losing preeminence in the supercomputer industry, it is important to examine key segments of the emerging computer/communications industry. Presently, three U.S. companies are selling the present version of the supercomputer for scientific and engineering computation--Control Data's ETA Systems, Cray Research and Denelcor. Two Japanese firms, Hitachi and Fujitsu, are also in this market. Fujitsu and Hitachi will deliver their first systems in the fourth quarter of 1983. All these American companies are delivering supercomputers and currently dominate the world market.

FIGURE 2

Examples of Convergence Between
Computer and Communication Technologies

1940-1990

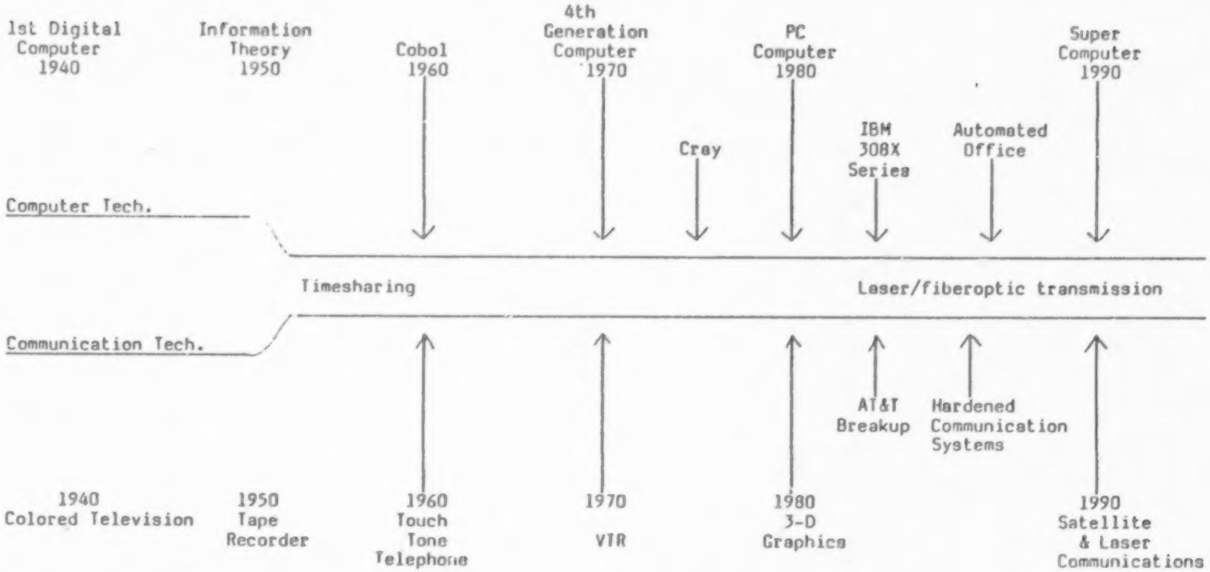
1st Digital Computer 1940
Information Theory 1950
Cobol 1960
4th Generation Computer 1970
Cray
PC Computer 1980
IBM 308X Series
Automated Office
Super Computer 1990
Computer Tech.
Timesharing
Laser/fiberoptic transmission
Communication Tech.
AT&T Breakup
Hardened Communication Systems
1940 Colored Television
1950 Tape Recorder
1960 Touch Tone Telephone
1970 VTR
1980 3-D Graphics
1990 Satellite & Laser Communications

Source: The Conference Board, "Information Technology," 1972, p. 192 and the Institute for Constructive Capitalism, The University of Texas at Austin.

The world market today is led by CDC and Cray Research. Fujitsu and Hitachi are preparing to deliver their current supercomputers, which they call "vector computers," in production quantities. One of the major attributes of their vector computers is that they can be made IBM compatible. Current IBM compatible machines can be tied to their vector computers. It is assumed that the price of the vector computer system will be equal to or less than the CDC, Cray or Denelcor units, none of which is currently IBM compatible.

Fujitsu's vector through-put computational capacity appears to be one-third to one-half higher than present American machines. Fujitsu has made arrangements for ICL, the major computer company in the U.K., and Seimens, the major computer company in West Germany, to market the vector machines in Europe for at least the non-IBM market. Current negotiations with a U.S. firm are underway to market the vector IBM compatible machine in North America. It is not clear who will market the IBM compatible vector machines in Europe or how they will be marketed. Fujitsu will market vector machines in Japan, Australia, Asia and Central and South America. It is believed that these machines will be non-IBM compatible. It is possible to benchmark applications on the Fujitsu vector machine to evaluate its IBM compatibility and to measure its other specifications. To the best of my knowledge, one successful, non-government American bench-marking application has already taken place.

Fujitsu's market in the U.S. will of necessity be earmarked for non-government related needs. Government stipulations generally require that over 50 percent of computer systems be American made. Because most supercomputers in America today are funded by the U.S. government through DOD, DOE, NASA and NSF, this market is the exclusive province of the three American firms.

The Fujitsu and Hitachi machines are reputed to have a larger computational capacity than current American counterparts. Most, if not all, Japanese and American companies have research underway to develop a giant 10,000 megaflop or more machine by the 1986-1990 time frame.

Denelcor is expecting to develop a 50 parallel processing supercomputer project. CDC has a 64 parallel processing supercomputer research project. Some universities are also researching a 50 or more parallel processor supercomputer. The competition is currently in the scientific area.

The marketing competition is for selected segments of the industry. As long as the supercomputer is sponsored primarily by the U.S. government, it will remain American. If it becomes non-government funded, the probability is that it can become largely Japanese. The critical dimension will be the source and availability of capital investments. The main issue that our scientific and advanced engineering communities face is how to secure the largest supercomputers for the least cost.

Scientific Competition--The previous section indicated that three U.S. firms and two Japanese firms are in a competitive mode for the supercomputer markets in the 1980's and 1990's. By this we simply mean that these five firms can deliver supercomputers to order today.

On the other hand, there are in both these nations significant research projects to develop new supercomputer technologies and equipments. It has been referred to in various media as "The Battle of the Supercomputers," "Japan's All-out Challenge to the U.S.," "The Race to Build a Supercomputer," "Japan's MITI Strike," "Japan Builds a Pillar of the Fifth Generation," "Japan--Here Comes Computer, Inc."

The scientific competition is not only between the U.S. and Japan. It is also a U.S. intrafirm competition; e.g., "Control Data and Cray Start Feuding Again Over Supercomputers." There are also competitive responses by the U.K. and the European community.

The worldwide "scientific competition" for supercomputers is presently underway. It is quite different from the earlier period during the 1950's and 1960's especially as to sources of funds, means by which research is administered and the manner of diffusing the innovations within the nation for participation in international markets. An overview by nation follows.

U.S.--Currently, supercomputers are funded by individual government agencies by ordering a supercomputer from one of the three U.S. firms. Each firm conducts its own directed research and development. Such efforts are generally funded through sales of company products. More specifically, their primary sources of funds are R & D expenditures that are tax deductible. New companies have a difficult time raising the necessary capital in the traditional financial markets because of the risk over 5-7 years of investing over $100 million to develop a first model. In addition, they do not have the necessary sales revenue and earnings for deduction of the R & D expenditures. Therefore, the probability is very low that as many new firms will emerge in the supercomputer arena as we saw in the 1950's and 1960's in the U.S. computer industry.

U.S. efforts in the scientific competition for the supercomputer industry can be summarized as follows:

1. The three U.S. firms have invested nearly 15 percent of their annual net revenues in supercomputer R & D for succeeding generations of computers.
2. Over twelve computer firms have agreed to a technological joint venture with a newly formed non-profit research oriented corporation--Microelectronics and Computer Technology Corporation (MCC). In the past four months, MCC has located in Austin, Texas with a working relationship with both flagship universities in Texas--The University of Texas at Austin and Texas A & M University.

   MCC is financed by the computer firms in the consortium through possible contributions for a ten-year period. Funds are pledged on a three-year cycle. Other firms are being encouraged to participate in the consortium. The participants will be permitted to use patents and other advances developed from in-house research and development. They will have a three-year lead time before such patents are made available to non-participating companies. Each participant can utilize the advances as it sees fit for its own product and service lines.

The MCC's four research areas are:

1. computer aided design and computer aided manufacturing (CAD-CAM);
2. software productivity;
3. advanced computer architecture; and
4. chip packaging.

These areas are applicable to the supercomputer as well as other computer/communication technology products and services.

It is expected that over the next ten years more than $1 billion could be invested in MCC research and development.

A current issue and concern for MCC involves antitrust matters. This issue may involve possible civil rather than Justice Department anti-competition suits based on the company's in-house research and proprietary control of the patents.

3. Another innovative type of R & D consortium is the Semiconductor Research Corporation (SRC) in Research Triangle Park, North Carolina. SRC was established in 1982 by the Semiconductor Industry Association as a subsidiary. SRC receives its funding from semiconductor firms, users of semiconductors and computer system firms. SRC in turn establishes centers of excellence at selected universities and gives contracts rather than grants to universities on a project basis. Currently, they plan to establish at least

eight university centers of excellence each funded at a minimum of $1 million per year. In addition, they have placed over 42 contracts with various universities at an average of $100 thousand per year. It is expected that over $100 million will be provided for research and development over the next decade.

The R & D thrusts of SRC are:

1. semiconductor automation;
2. semiconductor materials;
3. system and component interaction;
4. reliability;
5. quality assurance and tests;
6. small device structures;
7. research into new methods of device fabricators;
8. research into manufacturing processes; and
9. research into packaging and off chip interconnections.

SRC does not plan on proprietary control of patents and will not perform component R & D in-house. Patents and know-how will be made available through the various universities participating in the centers and projects. Through such efforts, SRC R & D can enhance both computer and communication technologies.

Both MCC and SRC have targeted the critical areas for advancing the state of the art for all computers including the supercomputer. Their programs clearly delineate what

U.S. computer and components industries believe to be the key bottlenecks facing most firms. Some companies are involved with both MCC and SRC. They are a distinct minority. Many of those not participating seem to be awaiting the resolution of the antitrust issues and other legal concerns.

4. Defense Advanced Research Projects Agency (DARPA). In March 1983, DARPA announced a new program which started as the "Nth Generation" but is currently known as Strategic Computing and Survivability (SCS) to develop a new generation of superintelligent computers for military use. The SCS program objectives are to develop a system which will allow it to see, reason, plan and supervise military systems in combat. DARPA's thrust will contain:

    1. Extension of artificial intelligence research;
    2. Multiprocessor architecture;
    3. Innovations in VLSI including:
        a. high speed and low power plus high degree of radiation hardness chips;
        b. gallium arsenide (GaAs);
        c. Pilot fabrication of GaAs;
        d. high performance, high density bulk memory for field application;
        e. new packaging including circuit technologies and interconnections (including electro-optics); and
        f. high performance bus structure and access to very high density storage.
    4. Rapid turnaround and fabrication of IC.

The SCS program will consist of joint projects being carried out between universities and industry at their respective locations as well as within the research laboratories of DOD. Close to $1 billion is being budgeted for this program over the next seven years.

The SCS program has targeted the following military applications:

1. Autonomous systems for advanced vehicles or satellite-control stations;
2. Collaborative systems to aid humans in controlling other machines;
3. Military assessment systems for war games, battlefield simulation, strategic and tactical alternative generation and solution, and combat operational data, command systems;
4. Expert advisory assistance for nuclear planning and logistics, flight operations and equipment maintenance.

The commercialization spin-offs of SCS will be developed under current regulations and governmental procedures.

5. Stanford University Center for Integrated Systems (SUCIS). SUCIS is the result of joint research sponsorship by the government, corporations and the University. The funding for the Center comes primarily from federal government grants (DOD, NIH, NSF, NASA and DOE). In addition, annual contributions from 19 corporations who are involved in the development and sale of computers and semiconductor components. They have already contributed $14.3 million for the design and construction of the Center, as well as for the research and design, fabrication, testing and applications of VLSI systems. The research and educational program consist of

about 60 faculty members, approximately a 150-member research staff and 300 Ph.D. candidates. Currently over $30 million is being spent annually.

Among the research tasks at Stanford are:

1. Knowledge based VLSI design;
2. VLSI information systems;
3. VSSI computer systems;
4. Medical and rehabilitation electronic sensors, circuits and systems;
5. Computer-aided, fast turnaround lab;
6. Integrated circuit process models;
7. Compound semiconductor and silica-on-insulator research;
8. Fundamental studies of semiconductor surfaces and interfaces.

Currently Stanford's SUCIS is concerned with the issues of academic freedom, patent rights, copyright privileges, royalty allocations and teaching resource dilution.

6. Microelectronic Center of North Carolina (MCNC). MCNC was created as a non-profit corporation in July 1980. It is funded primarily by the state of North Carolina which has agreed to provide approximately $43 million out of $50 million for 1981-1985. Participating academic institutions are: Duke, University of North Carolina at Chapel Hill, University of North Carolina at Charlotte, North Carolina State University, North Carolina A & T University and the Research Triangle Institute. Three corporations are

currently affiliated with MCNC.

The objectives of MCNC are:

1. Vertical integration of general and special-purpose machines needed for the next generation;
2. Translation of global basic research results into useful technology for industry and government.

The key projects include:

1. VLSI design systems (Vivid);
2. Fast prototyping facility for hybrid techniques;
3. Advanced silicon-wafer fabrication.

In addition, there are four projects that use vertical integrated capabilities. These are:

1. Pixel planes;
2. Boolean vector machine;
3. Functional program language machine; and
4. Image-processing computers.

Resultant royalties are divided and administered under each participating institution's policy.

Japan--The Japanese approach to supercomputers has taken a different strategy than that in the U.S. Their strategy is to identify and target markets. They have targeted for the 1990's two markets, one for scientific and engineering computation and a second for more intelligent processing of knowledge. The first is referred to as the Supercomputer Program, the second as the 5th Generation Computer Program.

Each program has the consent and support of the Japanese government. Two program teams consisting of university, firm and government researchers have been formed. Each is funded jointly by the government and six major firms in the computer industry (Fujitsu, Hitachi, NEC, Toshiba, Nippon Electric and NTT).

In broad terms, each of the supercomputer projects will be jointly funded for a ten-year period at $450-500 million dollars. Once a program is completed, the scientists and engineers will return to their respective institutions. In other words, they have built in what we call a "sunset" provision for the programs, which incorporates an automatic diffusion process for innovation developed during the term of the programs. This is quite a contrast to the diffusion process of the U.S. cooperative technology venturing of MCC and SRC.

The technologies being developed under the 5th Generation Program are in comparative infancy in relation to the program specifications. These "seed technologies" have been identified as follows:

1. very large scale integration;
2. high-speed devices;
3. communications;
4. parallel processing;
5. software; and
6. artificial intelligence and pattern recognition.

It is interesting to note that, even if the Japanese do not meet these designed specifications for the 5th Generation Computer System, they still expect to find many applications for other markets both before and after 1990. By advancing the boundaring of the technology's state of the art, the Japanese hope to finance a portion of the 5th Generation Program by commercializing earlier technological advances. While the exact nature of the by-products of the 5th Generation Program are difficult to predict, they will permit the Japanese to have an important presence in the following markets:

1. artificial intelligence;
2. computer-aided engineering;
3. decision support systems; and
4. intelligent robots.

The Japanese Supercomputer Program has two important scientifically competitive thrusts:

1. a lightning-fast supercomputer using Josephson junctions or high electron mobile transistors; and
2. an optical computer using optical fibers.

Most Japanese government funding for the supercomputer for scientific and engineering computation will take the form of 100 percent grants to industry. Work on these computers started in 1981. The government has loaned 30 researchers plus two engineers from each of nine Japanese firms to work in a laboratory loaned by Fujitsu. Their task is to develop the all gallium-arsenide devices needed for computing with light pulses in contrast to electronic signals. A major objective of this program is to develop the new components, including semiconductors, fiber optics and lasers. These components will in turn have commercial spin-off applications in computers and communications. They have targeted a supercomputer to be developed by 1990 with specifications a thousand times faster than the 1982 U.S. competitive computers. They expect that this will involve large scale, great speed parallel basic processing involving anywhere between 100 and 1,000 processors working simultaneously.

In addition, they expect to develop their own internationally competitive parallel processing software and hardware. This supercomputer program was started to meet the challenge of IBM's announcement in 1981 that it is planning to develop around 1985 a computer incorporating Josephson junction logic elements.

U.K.--For the first time in their history, the British have embarked on a collaborative research program among industry, government and academia that will achieve advances in four key sectors:

1. software engineering;
2. very large scale integration;
3. man-machine interfaces; and
4. intelligent knowledge-based systems.

This approach is a departure from past U.K. strategies. It is broadly modeled on the Japanese counterpart.

The cost of the program will be spread over five years and entail investments of $550 million with $316 million from government. Government funds will fund match 50 percent of all industrial funds. The work in the universities will be funded 100 percent by the government. A five-person directorate will supervise the program. The Secretary of the Science and Research Engineering will head the group while four industrial specialists (one from each of the key areas) will also serve on the directorate. The five-person directorate will report to a part-time board.

Europe--The European Commission is trying to establish a five-year European Strategic Program for Research and Information Technologies (ESPRIT). They are seeking funding from the participating countries by November 1983. The Esprit program plans to cover:

1. microelectronics;
2. softweare technology;
3. advanced information processing;
4. computer integrated manufacturing; and
5. office automation.

European Consortium--In the interim, three of the largest West European computer companies have formed a private joint research institute located in Bavaria. Compagnie (Machine), Bull (France), ICL plc (Britain) and Siemens AG (West Germany) will equally finance and run the Institute, which will focus on knowledge processing.

Summary

1. The scientific supercomputer competition has become a worldwide regional competition for the emerging global market. The current four active groups of competitors are U.S. firms, Japanese firms, the U.K. Program and the private transnational European research institute.
2. Two of the scientific competitors do not have major government funding--the U.S. companies and the European transnational institute. Defense R & D and VHSIC programs support the research of two competitors--the U.S. and the U.K.

3. Japan has extensive government support both in terms of funds and research scientists for advanced component research.
4. Spin-offs or by-products are being planned by two of the competitors--Japan and private U.S. firms.
5. While Japan tends to target markets, U.S. firms tend to target incremental technological break-throughs.
6. Only Japan has firm and dedicated funding for ten years. The U.S. firm funding is dependent upon the market and the individual firm's ability to generate earnings. (See Chart A).

CHART A*

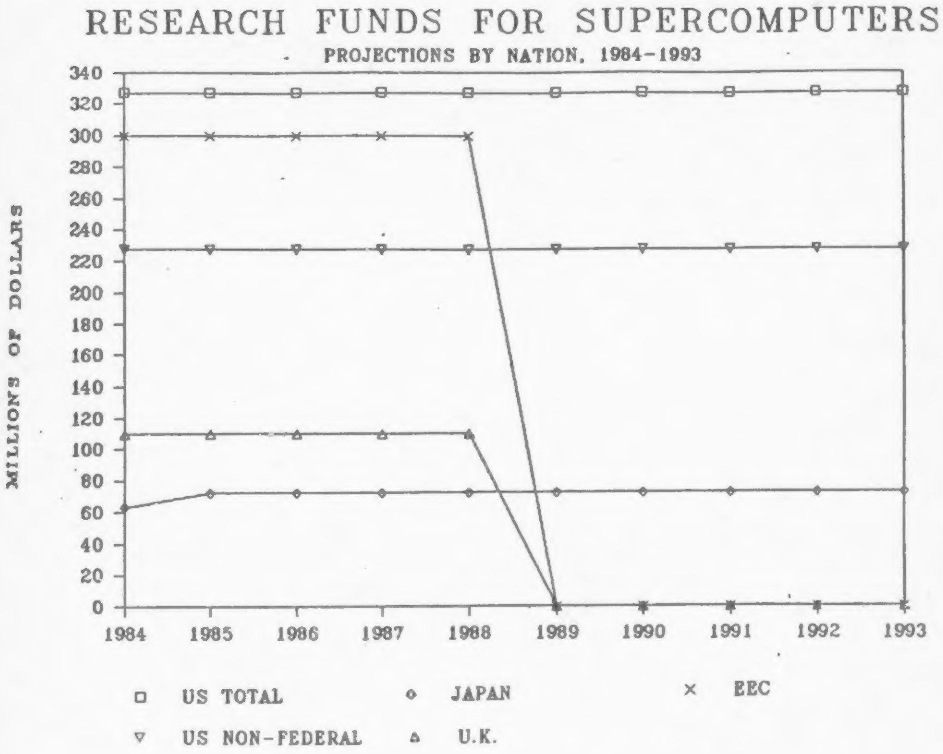

Source: Spectrum, November 1983, nad the Institute for Constructive Capitalism, November 1983.

* See following table for distribution of U. S. research funding.

CHART A

FORECAST
FUNDING DISTRIBUTION
U.S. SUPERCOMPUTER RESEARCH AND DEVELOPMENT
1984-1993

Research Institution	Estimated Ten-Year Expenditure in Millions of 1983 Dollars
Corporate Supercomputer R & D	$1180
Microelectronics and Computer Technology Corpotation (MCC)	1000
Semiconductor Research Corporation (SRC)	100
Stanford University Center for Integrated Systems (SUCIS)	300
Microeletronics Center of North Carolina (MCNC)	100
Subtotal-	$2300
Defense Advanced Projects Research Agency (DARPA)	1000
Total: U. S. Supercomputer Research	$3300

III. The Future of the Supercomputer Industry--A Forecast

If the supercomputer industry follows the same policies for preeminence as the computer industry did in the 1960's and 1970's, that is, if it builds a flourishing supercomputer industry rather than focus only on the use of the supercomputer and its technologies, then it is possible to forecast such results. Forecasting is an integral part of the innovation process; it is action and policy oriented. Its major assumption is that one can direct a path into the future which can be either self-fulfilling or self-defeating. Therefore, it is necessary to forecast supercomputer annual unit sales, their prices--high, medium and low--and the projected annual world supercomputer sales.

1. Annual Unit Supercomputer Sales Estimates.

Estimates in the current literature indicate that the market for supercomputers is mixed. There is little or no consensus concerning total sales of supercomputers. Fifty machines have been installed since the start of this year, and some predict a total market saturation point of 100 machines. Others project a maximum of 500 installed machines over the next decade. It is assumed for purposes of this testimony that these numbers are for non-IBM compatible operations. When IBM compatible projections are added, our projections indicate that the market could exceed 5,000 for the next decade. Chart B shows the annual supercomputer sales estimates.

CHART B

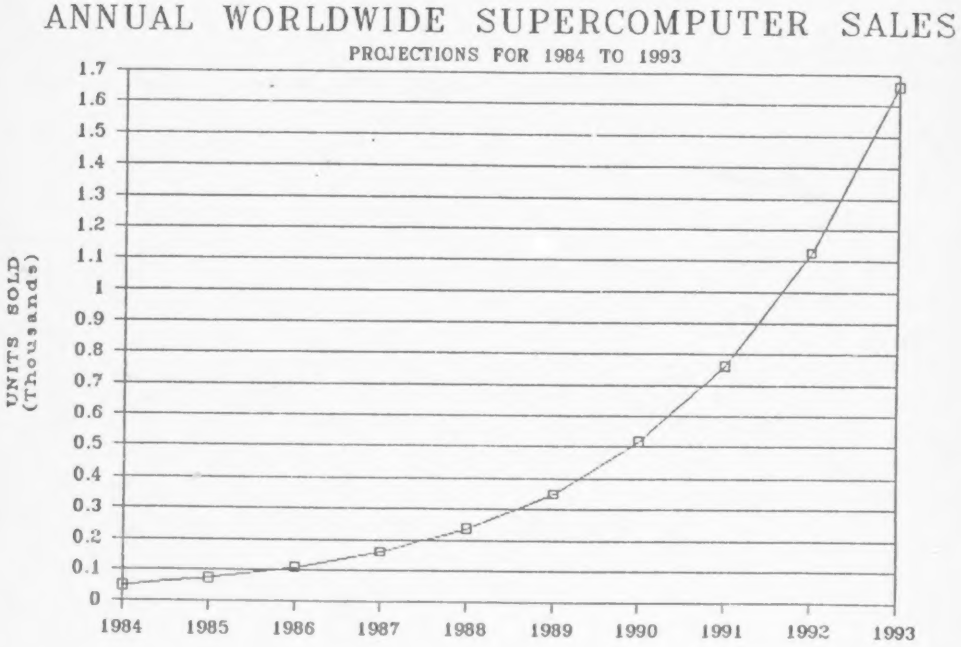

Source: The University of Texas at Austin, The Institute for Constructive Capitalism, November, 1983.

2. Price Projections.

In the computer systems industry, learning curve price projections are standard operating procedure. For purposes of estimating world supercomputer sales, Chart C provides three price projections using 90 percent, 85 percent and 80 percent learning curves. For the quantities shown in Chart B, the per unit prices are projected in Chart C. For comparative purposes, the years 1984 and 1993 are shown below:

	1984	1993
	(In Millions)	
90 percent learning curve	$10.0	$3.9
85 percent learning curve	$10.0	$2.3
80 percent learning curve	$10.0	$1.4

3. Projected Annual World Supercomputer Sales.

Chart D shows our estimates of annual world supercomputer sales. Predicted on quantities shown in Chart B and the price schedules in Chart C, our estimates of the total market annual sales of supercomputers are shown in Chart D. For comparative purposes, the total global sales for supercomputers by learning curves are shown below:

	1983	1993
	(In Billions)	
90 percent learning curve	$.5	$6.4
85 percent learning curve	$.5	$3.7
80 percent learning curve	$.5	$2.0

Our projections are significantly higher than others have made for the supercomputer market.

CHART C

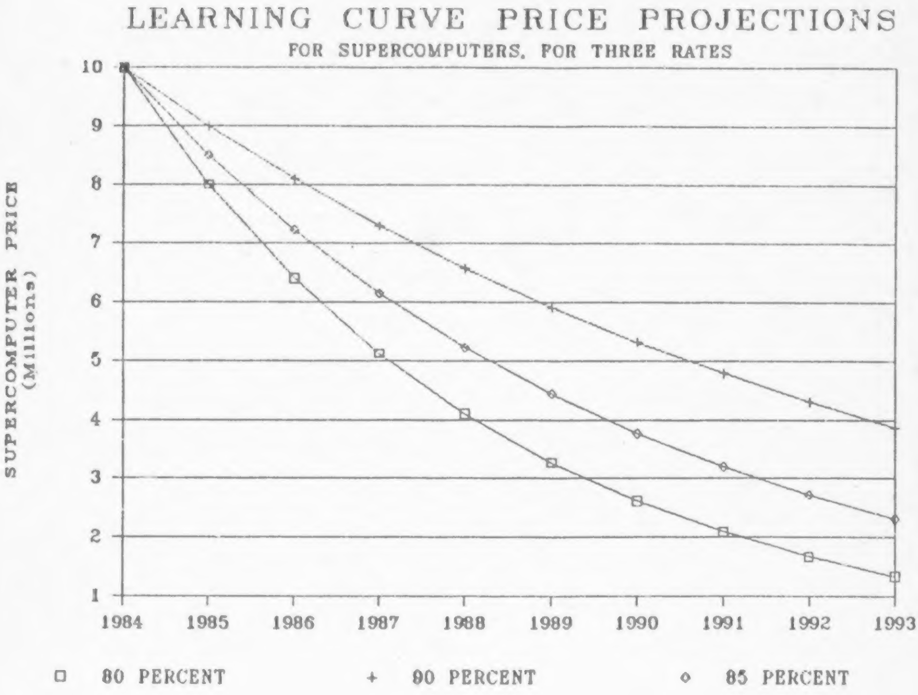

Source: The University of Texas at Austin, The Institute for Constructive Capitalism, November 1983.

CHART D

ANNUAL WORLDWIDE SUPERCOMPUTER SALES

PROJECTED REVENUE SENARIOS, 1984-1993

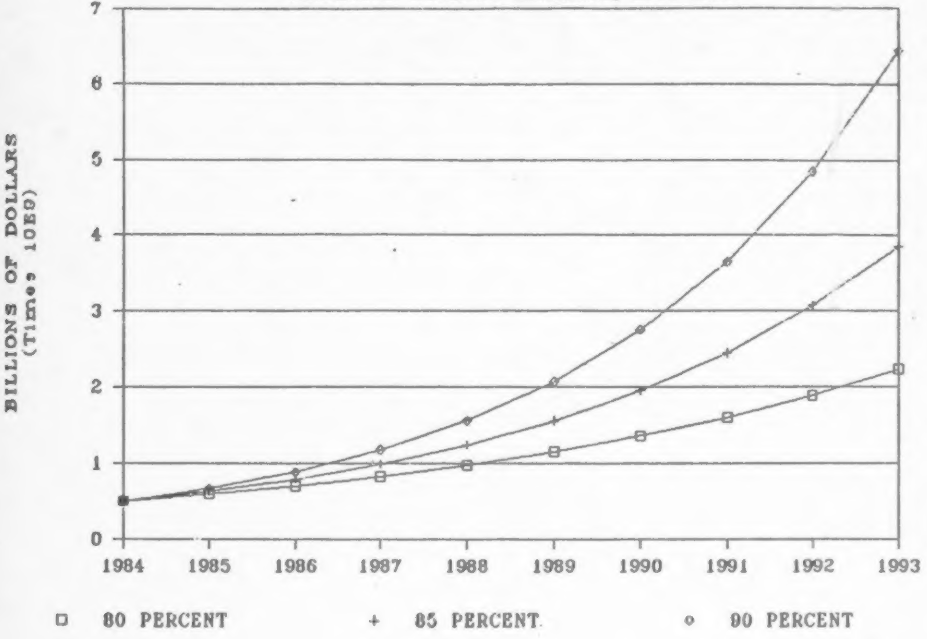

Source: The University of Texas at Austin, The Institute for Constructive Capitalism, November 1983.

Chart E shows that a projection of IBM's annual world revenues, exclusive of supercomputers, indicates it will be over $170 billion. We have estimated that the total world computer market will be between $320 billion and $430 billion. Therefore, our forecast for supercomputers is less than 2 percent of the 1993 world computer industry market. IC^2's forecast would have to be quintupled before it equaled 10 percent of the then forecasted world computer market.

On the other hand, high sales of $6.4 billion of supercomputers by 1993 could create between 80,000 and 107,000 new jobs. In addition, the industry could require between 100,000 to 250,000 new jobs for operating the supercomputers. These are extremely significant employment forecasts. This employment is about equal to today's total U.S. employment in industries such as furniture and fixtures, and miscellaneous manufacturing; or one and one-half times the employment in petroleum and coal products; or about one and one-half times the employment in leather and leather products. How the U.S. competitively shares the market is, therefore, very important to consider. Having the biggest, fastest and most expensive machine may not be enough in assessing economic consequences.

In order to emphasize the economic impact of losing preeminence in the marketplace for supercomputers, it might be helpful to review the recent 64K RAM history as an example.

The 64K RAM got off to an early start in Japan through a technology venture-sharing, government-sponsored program for research and development of very large scale integration in the 1979 timeframe.

CHART E

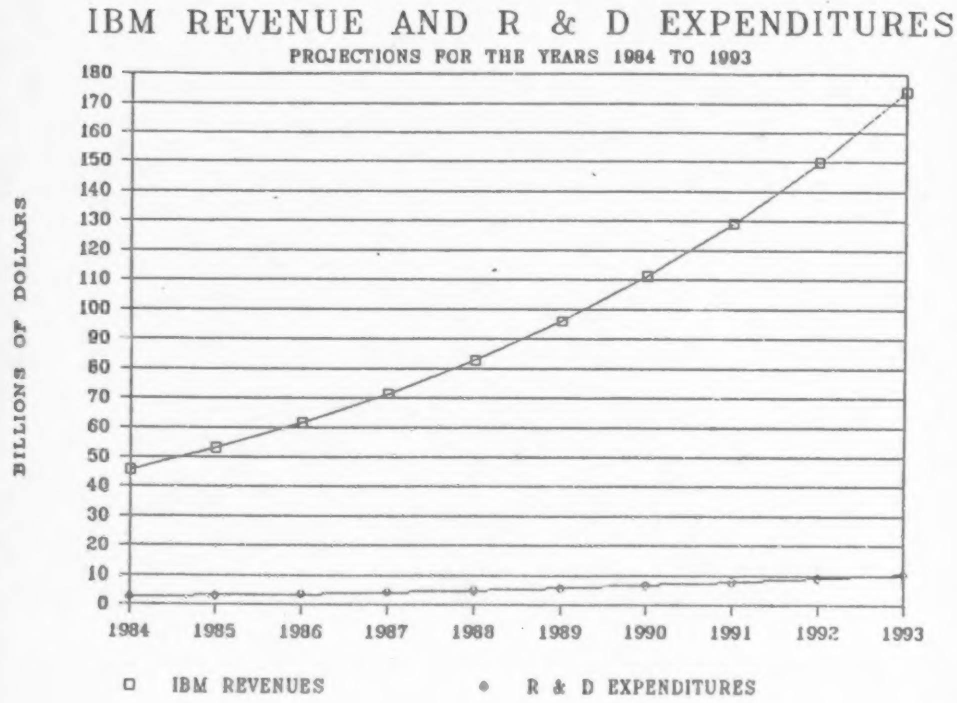

Source: First Boston, Insights, October 4, 1983, and The University of Texas at Austin, The Institute for Constructive Capitalism, November 1983.

The Japanese 64K RAM designs were targeted for Nippon Telegraph and Telephone Public Corp. equipment. The Japanese government-sponsored research included three Japanese companies--Fujitsu, NEC and Hitachi. By 1980, Fujitsu was commercially offering the 64K RAM by the thousand chips per month. By 1981, the Japanese had captured 70 percent of the 64K RAM market. The 64K RAM world market is expected to be over $1.5 billion in 1983 and over $2.0 billion by 1985. It is remarkable how fast one can move and take world market share. In the period that the U.S. 64K RAM manufacturers were losing market share, the number of U.S. companies producing 16K RAMs dropped from 12 to only six that could produce 64K RAMs competitively.

Supercomputer systems are an important factor in developing new chips. Loss of the supercomputer market could well lead to loss of future advanced chips markets for the U.S. One way of looking at the magnitude of such a loss is to examine what impacts the loss of 64K RAM chip market is to U.S. chip manufacturers.

In the following analysis, it is assumed that the Japanese chip manufacturers are able to obtain 50 percent of the U.S. chip market by 1993.

TABLE 2

Projections for Chips
1983 and 1993
(In Billions of Dollars)

	Estimated World Market Demand	U.S. Market Demand	Other Nations Market Demand
1983	$16	$ 7	$ 9
1993	$90	$45	$45

A more detailed forecast of the U.S. share of the market by source is shown in Table 3.

TABLE 3

U.S. Market Chip Demand
1983 and 1993
(In Billions of Dollars)

		Source of Supply for Chips		
	U.S. Market Demand	Internal System Mfg.	U.S. Chip Mfg.	Japanese Chip Mfg.
1983	$ 7	$2.3	$3.5	$1.2
1993	$45	$22.5	$11.25	$11.25

The forecast of chips for U.S. demands would increase from $7 billion to $45 billion or by more than 600 percent in the next decade. About one-half of the demand would be provided by system manufacturers own chip producing facilities. Over one-quarter would be provided by the Japanese and the U.S. chip manufacturers. The Japanese-U.S. chip market could increase ten fold in the coming decade. Meanwhile, the U.S. chip manufacturers would increase their domestic sales opportunity over three fold. The 64K RAM Japanese technical venture has provided a great opportunity for growth for Japanese chip manufacturers. This story of the 64K RAM when coupled to Japanese abilities to design supercomputers that are both IBM compatible and non-compatible provides a very important lesson when considering the future of the U.S. supercomputer industry.

When the total worldwide industry is dominated by IBM, it is important that U.S. supercomputer manufacturers including IBM be involved in the development, production and sales of supercomputers.

If Japanese manufacturers capture the whole of the IBM-compatible supercomputer market, they would have over 90 percent of our forecasted market. The remaining 10 percent would also be shared by the Japanese as their machines are designed for both the IBM compatible and non-compatible markets. Furthermore, if they had the bulk of production volumes, they could sell at much lower prices than U.S. manufacturers in the 1984-1993 time frame. (See Chart C).

Chart F forecasts supercomputer unit sales for the 1983-1993 time frame by U.S. and Japanese companies. The forecast assumes that IBM does not enter the supercomputer market until after 1993 nor do the other U.S. suppliers who could provide IBM compatible supercomputers. However, the U.S. suppliers provide 95 percent of the non-IBM compatible sales. The forecasted annual sales for 1983-1993 are shown in Chart G. For comparative purposes, sales for 1983 and 1993 are shown below:

	Forecast	
Annual Sales	Japanese Company Sales (In Billions of Dollars)	U.S. Firm Sales (In Billions of Dollars)
1983	$ .01	$ .5
1993	$2.50	$1.0

If these assumptions prevailed, the Japanese would clearly dominate the world supercomputer market during the decade. U.S. companies under present plans would not become economically competitive until they introduced their newer models which are currently under scientific competition with the Japanese.

CHART F

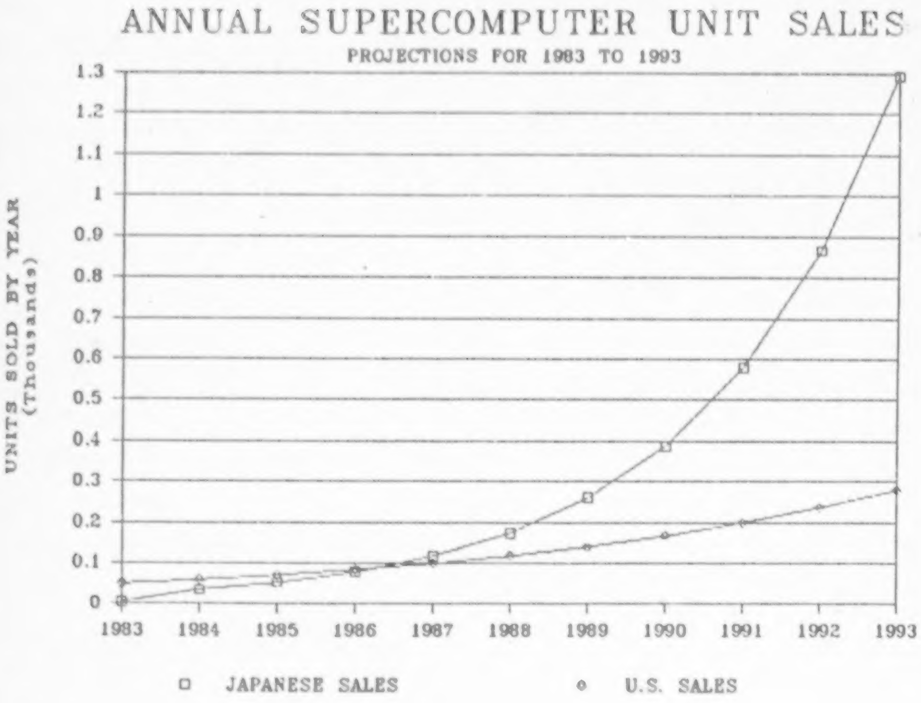

Source: The University of Texas at Austin, The Institute for Constructive Capitalism, November 1983.

CHART G

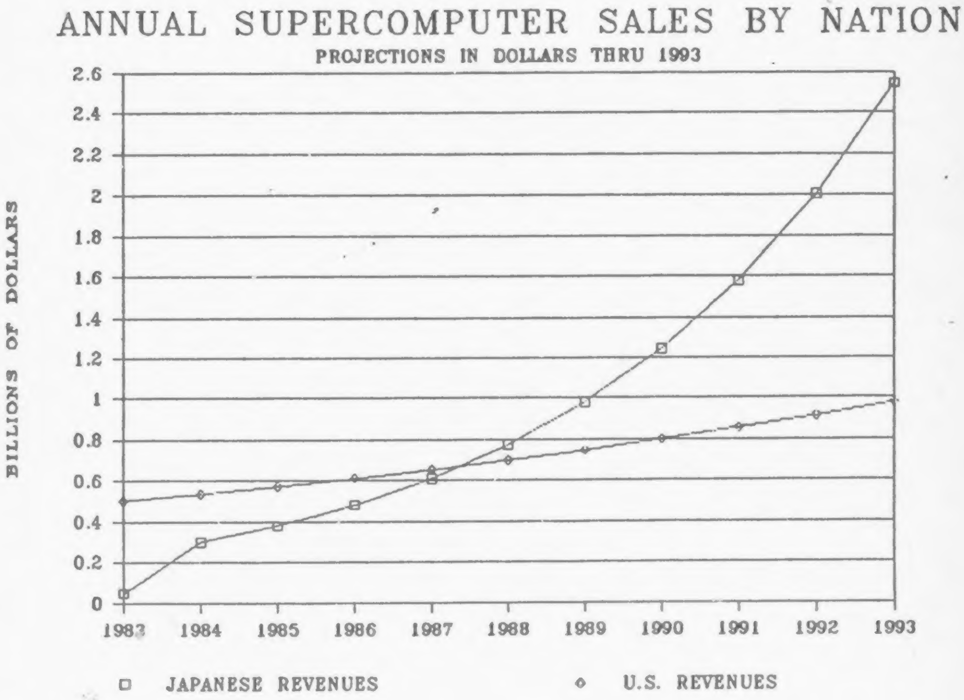

Source: The University of Texas at Austin, The Institute for Constructive Capitalism, November 1983.

The Japanese firms have already made commitments to the current supercomputer market. This means that they have the ability to meet market demands as shown in Chart B with relatively small capital investments. In other words, Japan is in the same position in 1983 to meet rapidly emerging markets for supercomputers as they were with the 64K RAM in 1980.

IV. Economic Consequences of Losing Preeminence in Supercomputers

The U.S. is engaged in an international scientific competition in the area of supercomputers. We are also engaged in an international scientific competition in telecommunications. Both of these scientific races are essential for what we have called the computer/communications industry for the 1990's and beyond.

The value of scientific exploration and discovery hardly needs to be emphasized to this distinguished committee. Scientific advances are basic to the continuation of our nation's economic growth. When appropriately commercialized, they can be used to reduce unemployment and underemployment, increase productivity, strengthen our national security, enhance our quality of life and improve our international balance of trade. But, scientific research is more than an economic good.

Scientific preeminence is a reflection of the American culture as well as its enterprise system. Most of us take great pride in the fact that all Nobel Prizes in science for 1983 were awarded to American citizens. Some of us still remember the bitter taste of loss of prestige when Sputnik was successfully launched by the U.S.S.R. in 1957. More importantly, the products and services from scientific achievements are marvelous attributes of our unique American culture. Other nations have the same desire. Consequently, the scientific race for the supercomputer must be viewed from the perspectives of both national prestige and economic well-being.

The U.S. can win the supercomputer scientific race but that does not ensure economic success in a global market. Other nations have become adept at applying American scientific advances to their own needs while improving their manufacturing, financial, marketing and managerial abilities to compete vigorously in international markets. Japan, the U.K. and perhaps the European community in the 1980's will focus more strongly on R & D and the innovation process in their supercomputer programs. So they may also be as effective in the scientific competition as the U.S. As a result, each nation may gain in prestige from its own unique scientific research, organizational R & D structures and massing of technological and human resources.

In this context, the ability to market a product a few months or a few years earlier may be of great value to an individual firm's well-being (and its immediate community). Current economic theory does not emphasize this critical time component. That is, free trade theory maintains that a rational person buys from the cheapest producer even if that means waiting for the product. Individual firms often indicate that such strategy can place us in economic jeopardy--i.e., loss of markets at home and abroad. Current economic theory has not advanced to the point where it can cope with rapid diffusion of scientific advances, extremely short technological life-cycles, ever increasing fixed capital investments with shorter economic lives than IRS regulations or tax incentives, comparatively brief generation times for new products and the rapid appearance and disappearance of markets. In summary, economic theory itself needs to be subjected to intense study to handle these newer classes of problems--more specifically, to measure firms' and nations' efficiency, effectiveness, flexibility and adaptability in a global market economy when competing with many sophisticated and highly-industrialized nations.

In this setting, there are a number of economic consequences that can be evaluated. First, the global computer industry during the 1980's can be progressively less dominated by U.S. companies. The 1990's will be even more competitive than the 1970-1980 era. This will not be the result of loss of scientific preeminence. Rather, it is a natural progression to lose share of global markets when other nations become directly competitive. The key question is what share of the future computer/communications market can/should the U.S. realistically expect to maintain? The global market for computer/communications is estimated to increase from $100-150 billion in 1983 to $500-700 billion, in constant 1982 dollars, in 1993. Currently, the United States scientific and technological preeminence has resulted in our estimation to be at least 75 percent of the current market or $75-120 billion. If U.S. preeminence in its selected fields continues, its share could drop to 50-60 percent. Therefore, its 1993 market would be between $250-420 billion, in constant 1982 dollars. Other competitive nations, including Japan, the U.K. and European common market nations, would increase market share from $25-30 billion to $250-280 billion under these assumptions. Economically, all would prosper even though scientific and technological preeminence was shared. In short, it is difficult to imagine the U.S. maintaining its current market share over the longer run.

The U.S. supercomputer industry share of worldwide markets could be dramatically below 50 percent by 1987-88 as shown in Chart D. IBM or others must enter the market for the U.S. to maintain 50 percent or more. IBM could well be ready to announce soon its "supercomputer"; but there is no overt evidence available on this. The loss of American market share could have a relatively immediate economic consequence in the form of price wars for computers and selected semiconductor chips during the next five years. This could result in less R & D funds for U.S. firms to continue developing the next generation of supercomputers.

Second, the supercomputer scientific competition makes clear that the volume of resources, including human talent, devoted to fundamental science and to technological innovations is limited. In simple supply and demand terms, this means that science and engineering human resources when scarce will be diverted to increasing returns on investments rather than on maintaining scientific preeminence. A useful way to make this point is illustrated by the State of California study which showed that there was a shortage of scientists and engineers to meet their estimated demand for both defense and space markets as well as non-defense markets.[3] The choice in pure economic theory would be to select those demands which provide the highest marginal yield. According to managerial theory, this means selecting those technological products and services that provide the highest gross margin as well as rate of return on investment commensurate with liquidity. Since the results of science do not enter directly into the private economy, the demand for prestige and academic excellence may require public support under such economic conditions.

---

[3] State of California, "The Effect of Increased Military Spending in California," Office of Economic Policy, Planning and Research, May 19, 1982.

To maintain scientific preeminence and a significant economic share of the supercomputer market, the U.S. government should consider ways of supporting the domestic market for American supercomputers. One way to support the domestic market while increasing the number of scientists, technicians and engineers is for the government to assist in placing supercomputers in major universities and non-profit research institutions across the country.

A more rapid dispersion of supercomputers than the market currently dictates will help to ensure the transfer of technology which in turn can support widespread continued R & D activities that will expand the whole computer/communications industry. At the same time, this dispersion will expand the knowledge and experience available to our human resource base. Supercomputers are still not available in many of our most prestigious academic institutions.

In addition, the supercomputer is needed for other scientific developments in key areas such as defense, space, astronomy, biology, meteorological forecasting, chemistry, genetics, biotechnology, national policy simulation and assessment and large-scale project management.

Third, the viability of the computer/communications industry depends on not only maintaining but also expanding private sector

R & D funding through revenues and earnings. If revenues drop 5-10 percent, our major computer firms will be in jeopardy and may be forced to cut their R & D activities, including work on the supercomputer. (See Chart H.)

The supercomputer for scientific computation has highlighted the need to recognize that the real issues do not involve catching up with the science or technology or even embracing the Japanese management system or government policies. Rather, they are:

1. Scientific achievements contribute to a nation's pride and heritage.

2. The most critical economic problem is centered on how to convert advances in science to technology resources, i.e., to put them economically into salable products that meet marketplace needs on a global basis. Having the fastest gallium-arsenide chip in the development laboratory does not assure tomorrow's supercomputer market. That is the lesson America has learned from the Japanese competition for the 64K RAM chips.

3. The ability to manufacture at lowest cost as well as meet very tight delivery dates are important economic considerations if preeminence in the marketplace is to be maintained. Advances in chip technology require very large capital investments. There is no end in sight for scientific advances in the density of chips, for the cost of capital or for the reduction in selling price of the unit (except in theory), as shown below:

CHART H

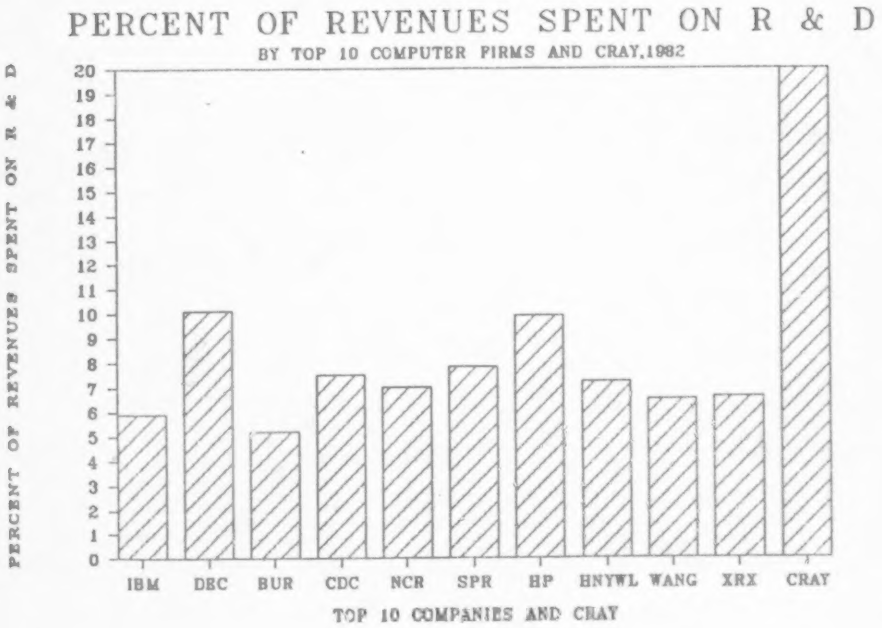

Source: Datamation, June 1983.

Year	# of Devices on Chip	Selling Price in Production	Est. Research Capital Invest. (In Mil. of $)	Refurbishing of Capital (In Years)
1960	50	$10.00	$ 5-10	5-10
1970	1,000	1.00	15-20	7
1980	500,000	.01	50-70	2-3
1990	20,000,000	.001	90-120	1-2

The ratio of sales revenues to capital investments is changing dramatically. In the 1970's, the ratio was $5.50 of capital investment to every $1 of revenue. In the 1980's, it is forecasted to be $2.00 per $1 of revenue. In the 1990's, the ratio could be a staggering $1.00 for $1.00.

We are already beginning to see a reduction of the number of firms in the chip industry. When 16K RAM chips were in production in the U.S., there were 12 firms. The 64K RAM chips saw the number reduced to three firms. There are currently four firms in the U.S. who say they will produce the 256K chips. Concentration in fewer firms may be competitively necessary for the U.S. to maintain preeminence in the international economy. It is our forecast that Japan will concentrate their supercomputer and large scale computers in two firms before 1990.

Chips are one of the key economic elements in determining a firm's ability to meet the world competition and in turn maintain U.S. economic preeminence. First, the architecture of the supercomputer is dependent on the chips characteristics. Chips, therefore, provide a competitive edge if they're faster or cheaper or do more functions or if they combine all three criteria. Second, the ability to make rapid and economical changes in the design stage, the manufacturing process or in the correction of latent defects found at the user's facilities will be of more critical importance in the future under global competition. The costs of changes grows at a compounded rate given geographic distance and progress in the life cycle of the computer generation.

These distance and progress issues compound the supercomputer firms' capital investment problem. They will have to provide much of their own chip production facilities not only to meet external demands but also to provide for their own internal needs. This is already true of the two major Japanese firms.

There are several economic consequences of such phenomena for national policy consideration. Recent forecasts indicate that major firms in the computer/communications industry will produce about one-half of the total chip market in 1983 or $45 billion for their own uses. They will buy the rest from chip manufacturers--domestic and foreign. Because of the nature of chip production processes, supercomputer manufacturers will be in a position to sell their excess capacity chips that do not meet their internal needs but could be used by others. Price competition, including dumping, could well become as common in chips as in steel today.

We are already seeing that our more creative chip manufacturers are unable to support their future capital needs for both fixed and working capital. Some, like Intel, have already sold a substantial equity position to IBM. Recently, Digital Equipment and Honeywell have made large equity investments in Trilogy reputedly for the right to the chip and new packaging developments. This could well lead to interesting merger and acquisitions as well as equity joint ventures. These could have implications for review of current security regulations, tax legislations and antitrust policies.

From a national technology perspective, the implications are that the current technology venturing arrangements with MCC and SRC may prove adequate for the scientific competition. However, the economic consequences of not being able to produce cheaply enough and deliver in required time may well result in production oriented technology ventures between medium and smaller sized supercomputer hardware and peripheral manufacturers. It might be necessary to consider the possibility that such technology ventures could include foreign firms. This in turn may necessitate a review of national policy issues of "Buy American," technology export and financing and extension of high-technology venture capital business.

4. Continually advancing supercomputer architecture, that is, the design of the computer system, is essential in maintaining economic preeminence. However, the process of designing succeeding generations of the supercomputer system is currently facing critical economic problems.

The period between generations of large scale computers, just as that between chips, has been substantially reduced. In the 1960's, a generation was 10 years; by the 1970's, the generation was about seven years. Currently, the generation has shrunk to 3-4 years. Put simply, this means that a firm's R & D personnel must design two new computer systems in the 1980's in the same time that it took to design one in the 1970's. In the words of one CEO with whom I talked, "There is no way that we can do that. It is impossible." What makes this situation worse is the current shortage of skilled personnel. This makes it difficult to keep up with competition domestically and internationally. During the past several years, the salaries and other incentives for key technical talent has risen dramatically. "Frustration" is the only word that describes the current state of our supercomputer manpower situation. The alternative to educating more computer scientists, engineers and technicians is a longer-term solution when the economic consequences are taken into account. For example, now that the Japanese have 70 percent of the U.S. 64K RAM chip market, they expect to maintain at least a 50 percent share of U.S. 256K RAM chip market. Shortage of personnel could well be a serious constraint to increasing or even maintaining the U.S. share of the global computer/communications industry market. It certainly permits foreign competition to target on gaps in our large domestic market and the rest of the world.

National policy must focus on how to continue the U.S. as a superpower in the computer and communications markets. The supercomputer is the single greatest impact on world communication, automated factories, health care delivery, biotechnology production, renewal of basic industry and heightened productivity of the service industry, including goverment. The real task is how to develop appropriately integrated policies, regulations and support mechanisms that extend the U.S. computer/communications industry. The commercialization of supercomputers for the global market is so tightly structured from scientific exploration to ultimate use, regeneration time is so short, investments so large and risk so great, that we cannot leave policy considerations to evolve accidently and independently as in the past. One needs only to look at the automobile industry to realize the possible consequences of short-sighted policy development. The computer/communications industry is our future source of large scale employment; i.e., one out of six. Table 3 emphasizes the importance of the supercomputer as a driver for the computer/communications industry.

TABLE 3

Effects of Preeminence in Supercomputer Technology on the Computer/Communications Industry

Drivers	Area Effected
Industries Developed and Expanded by Super Computer Related Technologies	Robotics/Industrial Automation
	Computer-aided Engineering
	Computer-aided Design
	Computer-aided Manfacturing
	Computer-aided Testing
	Computer-aided Quality Control
	Biotechnology
	Office Automation
	International Financial Services
	Service Industry
	Computer Industry
	Semiconductor Industry
Revitalization of Basic Industries	Energy
	Petroleum
	Nuclear
	Steel
	Automobile
	Textile
	Chemicals
Institutions Receiving Significant Productivity Boosts	Government
	Defense
	Education
	Research
Communications Technology Enhancements	Subscribers Equipment
	Telephone Co. Equipment
	Overseas Carriers
	Satellite Carriers

The consequences of losing economic and scientific preeminence in the supercomputer industry are vast. The supercomputer is a central driver for the rapidly emerging worldwide computer/communications industry. It impacts communication developments, the renewal of basic industries, productivity increases and the development and expansion of new industries. It is essential in improving our educational structure, fulfilling critical manpower requirements and enhancing our industrial creativity and innovation. It is the seed for encouraging the emergence of a myriad of technology venture businesses in the context of a private enterprise system that has always been the unique American way to achieve and maintain U.S. economic and scientific preeminence.

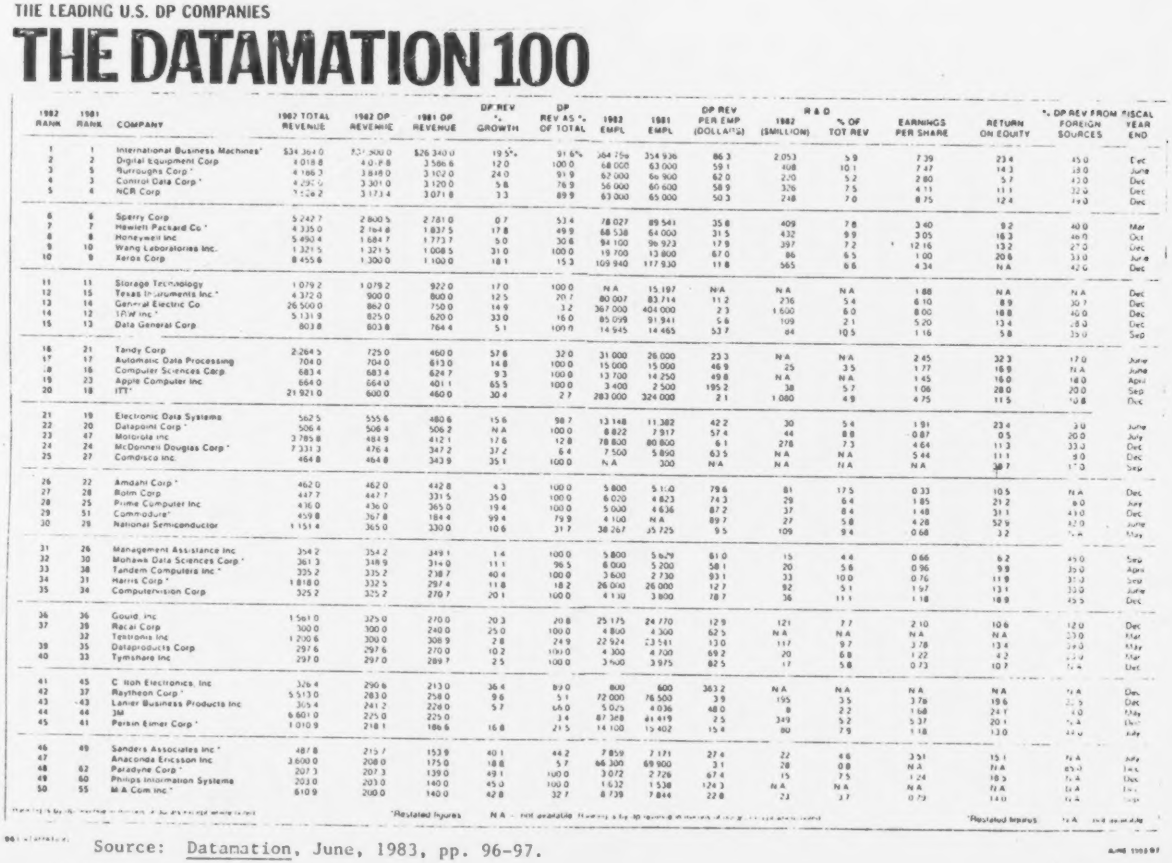

THE LEADING U.S. DP COMPANIES

THE DATAMATION 100

1982 RANK	1981 RANK	COMPANY	1982 TOTAL REVENUE	1982 DP REVENUE	1981 DP REVENUE	DP REV % GROWTH	DP REV AS % OF TOTAL	1982 EMPL	1981 EMPL	DP REV PER EMP (DOLLARS)	R & D 1982 ($MILLION)	R & D % OF TOT REV	EARNINGS PER SHARE	RETURN ON EQUITY	% DP REV FROM FOREIGN SOURCES	FISCAL YEAR END
1	1	International Business Machines*	$34 364.0	$31 500.0	$26 340.0	19.5%	91.6%	364 796	354 936	86.3	2 053	5.9	7.39	23.4	45.0	Dec
2	2	Digital Equipment Corp	4 018.8	4 018.8	3 586.6	12.0	100.0	68 000	63 000	59.1	408	10.1	7.47	14.3	38.0	June
3	5	Burroughs Corp*	4 186.3	3 848.0	3 102.0	24.0	91.9	62 000	66 900	62.0	220	5.2	2.80	5.7	43.0	Dec
4	3	Control Data Corp*	4 29[illegible]	3 301.0	3 120.0	5.8	76.9	56 000	60 600	58.9	326	7.5	4.11	11.1	32.0	Dec
5	4	NCR Corp	3 526.2	3 173.4	3 071.8	3.3	89.9	63 000	65 000	50.3	248	7.0	8.75	12.4	49.0	Dec
6	6	Sperry Corp	5 242.7	2 800.5	2 781.0	0.7	53.4	78 027	89 541	35.8	409	7.8	3.40	9.2	40.0	Mar
7	7	Hewlett Packard Co*	4 335.0	2 164.8	1 837.5	17.8	49.9	68 538	64 000	31.5	432	9.9	3.05	16.3	46.0	Oct
8	8	Honeywell Inc	5 490.4	1 684.7	1 773.7	5.0	30.6	94 100	96 923	17.9	397	7.2	12.16	13.2	27.0	Dec
9	10	Wang Laboratories Inc.	1 321.5	1 321.5	1 008.5	31.0	100.0	19 700	13 800	67.0	86	6.5	1.00	20.6	33.0	June
10	9	Xerox Corp	8 455.6	1 300.0	1 100.0	18.1	15.3	109 940	117 930	11.8	565	6.6	4.34	N A	42.0	Dec
11	11	Storage Technology	1 079.2	1 079.2	922.0	17.0	100.0	N A	15 197	N/A	N A	N A	1.88	N A	N A	Dec
12	15	Texas Instruments Inc*	4 372.0	900.0	800.0	12.5	20.7	80 007	83 714	11.2	236	5.4	6.10	8.9	30.7	Dec
13	14	General Electric Co	26 500.0	862.0	750.0	14.9	3.2	367 000	404 000	2.3	1 600	6.0	8.00	18.8	40.0	Dec
14	12	TRW Inc*	5 131.9	825.0	620.0	33.0	16.0	85 099	91 941	5.6	109	2.1	5.20	13.4	28.0	Dec
15	13	Data General Corp	803.8	803.8	764.4	5.1	100.0	14 945	14 465	53.7	84	10.5	1.16	5.8	35.0	Sep
16	21	Tandy Corp	2 264.5	725.0	460.0	57.6	32.0	31 000	26 000	23.3	N A	N A	2.45	32.3	17.0	June
17	17	Automatic Data Processing	704.0	704.0	613.0	14.8	100.0	15 000	15 000	46.9	25	3.5	1.77	16.9	N A	June
18	16	Computer Sciences Corp	683.4	683.4	624.7	9.3	100.0	13 700	14 250	49.8	N A	N A	1.45	16.0	18.0	April
19	23	Apple Computer Inc	664.0	664.0	401.1	65.5	100.0	3 400	2 500	195.2	38	5.7	1.06	28.0	20.0	Sep
20	18	ITT*	21 921.0	600.0	460.0	30.4	2.7	283 000	324 000	2.1	1 080	4.9	4.75	11.5	10.8	Dec
21	19	Electronic Data Systems	562.5	555.6	480.6	15.6	98.7	13 148	11 382	42.2	30	5.4	1.91	23.4	3.0	June
22	20	Datapoint Corp*	506.4	506.4	506.2	N A	100.0	8 822	7 917	57.4	44	8.8	-0.87	0.5	20.0	July
23	47	Motorola Inc	3 785.8	484.9	412.1	17.6	12.8	78 800	80 800	6.1	278	7.3	4.64	11.3	33.0	Dec
24	24	McDonnell Douglas Corp*	7 331.3	476.4	347.2	37.2	6.4	7 500	5 890	63.5	N A	N A	5.44	11.1	9.0	Dec
25	27	Comdisco Inc.	464.8	464.8	343.9	35.1	100.0	N A	300	N/A	N A	N A	N A	38.7	17.0	Sep
26	22	Amdahl Corp*	462.0	462.0	442.8	4.3	100.0	5 800	5 100	79.6	81	17.5	0.33	10.5	N A	Dec
27	28	Rolm Corp	447.7	447.7	331.5	35.0	100.0	6 020	4 823	74.3	29	6.4	1.85	21.2	8.0	July
28	25	Prime Computer Inc	436.0	436.0	365.0	19.4	100.0	5 000	4 636	87.2	37	8.4	1.48	31.1	41.0	Dec
29	51	Commodore*	459.8	367.8	184.4	99.4	79.9	4 100	N A	89.7	27	5.8	4.28	52.9	42.0	June
30	29	National Semiconductor	1 151.4	365.0	330.0	10.6	31.7	38 267	35 725	9.5	109	9.4	0.68	3.2	N A	May
31	26	Management Assistance Inc	354.2	354.2	349.1	1.4	100.0	5 800	5 629	61.0	15	4.4	0.66	6.2	45.0	Sep
32	30	Mohawk Data Sciences Corp*	361.3	348.9	314.0	11.1	96.5	6 000	5 200	58.1	20	5.6	0.96	9.9	35.0	April
33	38	Tandem Computers Inc*	335.2	335.2	238.7	40.4	100.0	3 600	2 730	93.1	33	10.0	0.76	11.9	31.0	Sep
34	31	Harris Corp*	1 818.0	332.5	297.4	11.8	18.2	26 000	26 000	12.7	92	5.1	1.97	13.1	33.0	June
35	34	Computervision Corp	325.2	325.2	270.7	20.1	100.0	4 130	3 800	78.7	36	11.1	1.18	18.9	45.5	Dec
36	36	Gould Inc	1 561.0	325.0	270.0	20.3	20.8	25 175	24 770	12.9	121	7.7	2.10	10.6	12.0	Dec
37	39	Racal Corp	300.0	300.0	240.0	25.0	100.0	4 800	4 300	62.5	N A	N A	N A	N A	33.0	Mar
	32	Tektronix Inc	1 200.6	300.0	308.9	2.8	24.9	22 924	23 541	13.0	117	9.7	3.78	13.4	39.0	May
39	35	Dataproducts Corp	297.6	297.6	270.0	10.2	100.0	4 300	4 700	69.2	20	6.8	1.22	4.2	23.0	Mar
40	33	Tymshare Inc	297.0	297.0	289.7	2.5	100.0	3 600	3 975	82.5	17	5.8	0.73	10.7	N A	Dec
41	45	C Itoh Electronics, Inc	326.4	290.6	213.0	36.4	89.0	800	600	363.2	N A	N A	N A	N A	N A	Dec
42	37	Raytheon Corp*	5 513.0	283.0	258.0	9.6	5.1	72 000	76 500	3.9	195	3.5	3.78	19.6	22.5	Dec
43	43	Lanier Business Products Inc	365.4	241.2	228.0	5.7	66.0	5 025	4 036	48.0	8	2.2	1.68	24.1	4.0	May
44	44	3M	6 601.0	225.0	225.0		3.4	87 388	81 419	2.5	349	5.2	5.37	20.1	N A	Dec
45	41	Perkin Elmer Corp*	1 010.9	218.1	186.6	16.8	21.5	14 100	15 402	15.4	80	7.9	1.18	13.0	44.0	July
46	49	Sanders Associates Inc*	487.8	215.7	153.9	40.1	44.2	7 859	7 171	27.4	22	4.6	3.51	15.1	N A	July
47		Anaconda Ericsson Inc	3 600.0	208.0	175.0	18.8	5.7	66 300	69 900	3.1	28	0.8	N A	N A	85.0	Dec
48	62	Paradyne Corp*	207.3	207.3	139.0	49.1	100.0	3 072	2 726	67.4	15	7.5	1.24	18.5	N A	Dec
49	60	Philips Information Systems	203.0	203.0	140.0	45.0	100.0	1 632	1 538	124.3	N A	N A	N A	N A	N A	Dec
50	55	M A Com Inc*	610.9	200.0	140.0	42.8	32.7	8 739	7 844	22.8	23	3.7	0.79	14.0	N A	Sep

[illegible] *Restated figures N A — not available [illegible]

Source: Datamation, June, 1983, pp. 96-97.

THE LEADING U.S. DP COMPANIES

THE DATAMATION 100

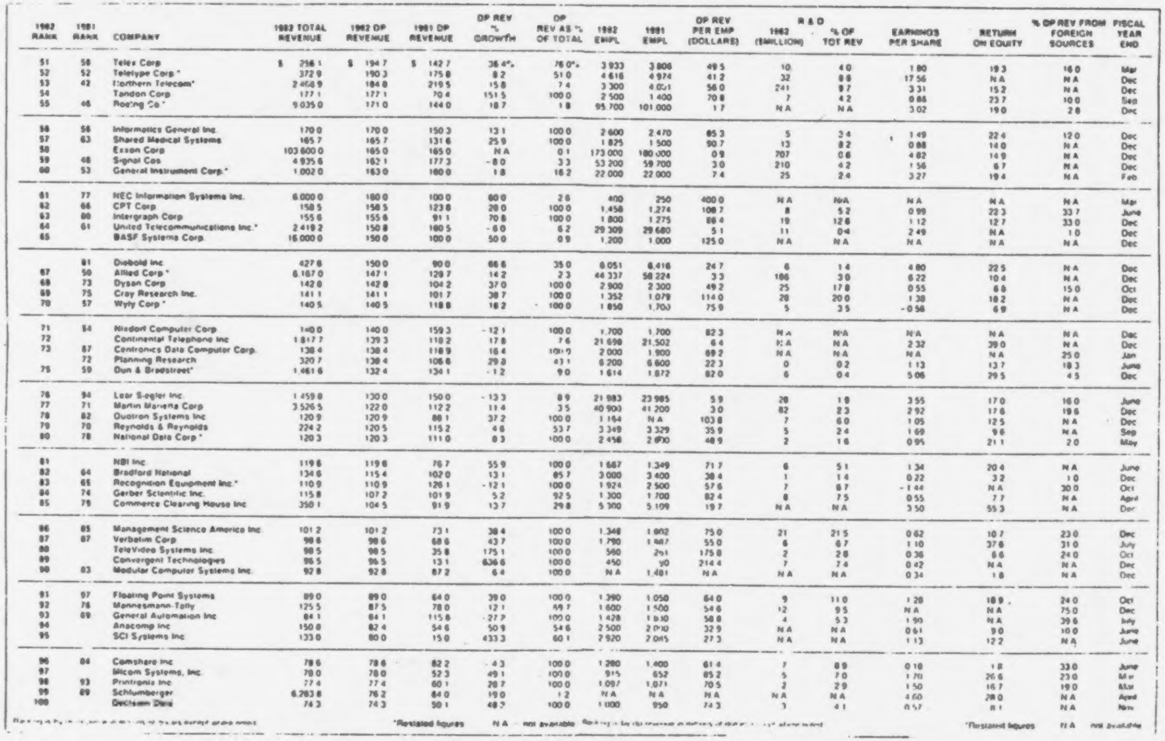

1982 RANK	1981 RANK	COMPANY	1982 TOTAL REVENUE	1982 DP REVENUE	1981 DP REVENUE	DP REV % GROWTH	DP REV AS % OF TOTAL	1982 EMPL	1981 EMPL	DP REV PER EMP (DOLLARS)	R & D 1982 ($MILLION)	R & D % OF TOT REV	EARNINGS PER SHARE	RETURN ON EQUITY	% DP REV FROM FOREIGN SOURCES	FISCAL YEAR END
51	58	Telex Corp	$ 256.1	$ 194.7	$ 142.7	36.4%	76.0%	3,933	3,806	49.5	10	4.0	1.80	19.3	16.0	Mar
52	52	Teletype Corp.*	372.9	190.3	175.8	8.2	51.0	4,616	4,974	41.2	32	8.8	17.56	N A	N A	Dec
53	42	Northern Telecom*	2,468.9	184.8	219.5	-15.8	7.4	3,300	4,051	56.0	241	9.7	3.31	15.2	N A	Dec
54		Tandon Corp	177.1	177.1	70.4	151.5	100.0	2,500	1,400	70.8	7	4.2	0.88	23.7	10.0	Sep
55	46	Boeing Co.*	9,035.0	171.0	144.0	18.7	1.8	95,700	101,000	1.7	N A	N A	3.02	19.0	2.8	Dec
56	56	Informatics General Inc.	170.0	170.0	150.3	13.1	100.0	2,600	2,470	65.3	5	3.4	1.49	22.4	12.0	Dec
57	63	Shared Medical Systems	165.7	165.7	131.6	25.9	100.0	1,825	1,500	90.7	13	8.2	0.88	14.0	N A	Dec
58		Exxon Corp	103,600.0	165.0	165.0	N A	0.1	173,000	180,000	0.9	707	0.6	4.82	14.9	N A	Dec
59	48	Signal Cos	4,935.6	162.1	177.3	-8.0	3.3	53,200	59,700	3.0	210	4.2	1.56	6.7	N A	Dec
60	53	General Instrument Corp.*	1,002.0	163.0	160.0	1.8	16.2	22,000	22,000	7.4	25	2.4	3.27	19.4	N A	Feb
61	77	NEC Information Systems Inc.	6,000.0	160.0	100.0	60.0	2.6	400	250	400.0	N A	N/A	N A	N A	N A	Mar
62	66	CPT Corp	158.5	158.5	123.8	28.0	100.0	1,458	1,274	108.7	8	5.2	0.99	22.3	33.7	June
63	80	Intergraph Corp	155.6	155.6	91.1	70.8	100.0	1,800	1,275	86.4	19	12.8	1.12	12.7	33.0	Dec
64	61	United Telecommunications Inc.*	2,419.2	150.8	160.5	-6.0	6.2	29,309	29,680	5.1	11	0.4	2.49	N A	1.0	Dec
65		BASF Systems Corp.	16,000.0	150.0	100.0	50.0	0.9	1,200	1,000	125.0	N A	N A	N A	N A	N A	Dec
	81	Diebold Inc.	427.8	150.0	90.0	66.6	35.0	8,051	8,416	24.7	8	1.4	4.80	22.5	N A	Dec
67	50	Allied Corp.*	6,167.0	147.1	128.7	14.2	2.3	44,337	58,224	3.3	186	3.0	6.22	10.4	N A	Dec
68	73	Dysan Corp	142.8	142.8	104.2	37.0	100.0	2,900	2,300	49.2	25	17.8	0.55	6.8	15.0	Oct
69	75	Cray Research Inc.	141.1	141.1	101.7	38.7	100.0	1,352	1,079	114.0	28	20.0	1.38	18.2	N A	Dec
70	57	Wyly Corp.*	140.5	140.5	118.8	18.2	100.0	1,850	1,700	75.9	5	3.5	-0.56	6.9	N A	Dec
71	54	Nixdorf Computer Corp	140.0	140.0	159.3	-12.1	100.0	1,700	1,700	82.3	N A	N/A	N/A	N A	N A	Dec
72		Continental Telephone Inc	1,817.7	139.3	118.2	17.8	7.6	21,690	21,502	6.4	N A	N A	2.32	39.0	N A	Dec
73	67	Centronics Data Computer Corp.	138.4	138.4	118.9	16.4	100.0	2,000	1,900	69.2	N A	N A	N A	N A	25.0	Jan
	72	Planning Research	320.7	130.4	106.8	29.3	41.1	6,200	6,600	22.3	0	0.2	1.13	13.7	18.3	June
75	59	Dun & Bradstreet*	1,461.6	132.4	134.1	-1.2	9.0	1,614	1,872	82.0	6	0.4	5.06	29.5	4.5	Dec
76	94	Lear Siegler Inc.	1,459.8	130.0	150.0	-13.3	8.9	21,983	23,985	5.9	28	1.8	3.55	17.0	16.0	June
77	71	Martin Marietta Corp	3,526.5	122.0	112.2	11.4	3.5	40,900	41,200	3.0	82	2.3	2.92	17.6	19.6	Dec
78	82	Quotron Systems Inc	120.9	120.9	88.1	37.2	100.0	1,164	N A	103.8	7	6.0	1.05	12.5	N A	Dec
79	70	Reynolds & Reynolds	224.2	120.5	115.2	4.6	53.7	3,349	3,329	35.9	5	2.4	1.69	9.6	N A	Sep
80	78	National Data Corp.*	120.3	120.3	111.0	8.3	100.0	2,458	2,690	48.9	2	1.6	0.95	21.1	2.0	May
81		NBI Inc.	119.6	119.6	76.7	55.9	100.0	1,687	1,349	71.7	6	5.1	1.34	20.4	N A	June
82	64	Bradford National	134.6	115.4	102.0	13.1	85.7	3,000	3,400	38.4	1	1.4	0.22	3.2	1.0	Dec
83	65	Recognition Equipment Inc.*	110.9	110.9	126.1	-12.1	100.0	1,924	2,500	57.6	7	6.7	-1.44	N A	30.0	Oct
84	74	Gerber Scientific Inc.	115.8	107.2	101.9	5.2	92.5	1,300	1,700	82.4	8	7.5	0.55	7.7	N A	April
85	79	Commerce Clearing House Inc	350.1	104.5	91.9	13.7	29.8	5,300	5,109	19.7	N A	N A	3.50	55.3	N A	Dec
86	85	Management Science America Inc.	101.2	101.2	73.1	38.4	100.0	1,348	1,002	75.0	21	21.5	0.62	10.7	23.0	Dec
87	87	Verbatim Corp	98.6	98.6	68.6	43.7	100.0	1,790	1,487	55.0	6	6.7	1.10	37.8	31.0	July
88		TeleVideo Systems Inc	98.5	98.5	35.8	175.1	100.0	560	251	175.8	2	2.8	0.36	6.6	24.0	Oct
89		Convergent Technologies	96.5	96.5	13.1	636.6	100.0	450	90	214.4	7	7.4	0.42	N A	N A	Dec
90	83	Modular Computer Systems Inc.	92.8	92.8	87.2	6.4	100.0	N A	1,481	N A	N A	N A	0.34	1.8	N A	Dec
91	97	Floating Point Systems	89.0	89.0	64.0	39.0	100.0	1,390	1,050	64.0	9	11.0	1.28	18.9	24.0	Oct
92	76	Mannesmann-Tally	125.5	87.5	78.0	12.1	69.7	1,600	1,500	54.6	12	9.5	N A	N A	75.0	Dec
93	69	General Automation Inc	84.1	84.1	115.6	-27.2	100.0	1,428	1,610	58.8	4	5.3	1.99	N A	39.6	July
94		Anacomp Inc	150.0	82.4	54.6	50.9	54.6	2,500	2,010	32.9	N A	N A	0.61	9.0	10.0	June
95		SCI Systems Inc	133.0	80.0	15.0	433.3	60.1	2,920	2,045	27.3	N A	N A	1.13	12.2	N A	June
96	84	Comshare Inc	78.6	78.6	82.2	-4.3	100.0	1,280	1,400	61.4	7	8.9	0.18	1.8	33.0	June
97		Micom Systems, Inc.	78.0	78.0	52.3	49.1	100.0	915	652	85.2	5	7.0	1.70	26.6	23.0	Mar
98	93	Printronix Inc	77.4	77.4	60.1	28.7	100.0	1,097	1,071	70.5	2	2.9	1.50	16.7	19.0	Mar
99	89	Schlumberger	6,283.8	76.2	64.0	19.0	1.2	N A	N A	N A	N A	N A	4.60	28.0	N A	April
100		Decision Data	74.3	74.3	50.1	48.3	100.0	1,000	950	74.3	3	4.1	0.57	8.1	N A	Nov

*Restated figures N A — not available

Source: Datamation, June, 1983, pp. 98-99.

Bibliography

"A Big Fight Over Tiny Chips," Time, August 9, 1982, p. 144.

"A Fifth Generation: Computers That Think," Business Week, December 14, 1981, p. 94.

Abegglen, James C. and Akio Etori, "Japanese Technology Today," Advertisement, Scientific America, November 1983, p. J4.

Ashley, Steven, "Keeping up With the Japanese: Defense Department to Push for Fifth Generation Computer," American Metal Market/Metalworking News, December 6, 1982, p. 7.

Banks, Howard and Kathleen K. Wagner, "A Matter of Life and Death," Forbes, February 28, 1983, p. 35.

Barnes, Clifford, "Who Will Pay for Supercomputer R & D?" Electronics, September 8, 1983, p. 87.

Beresford, Roderic, "Opening Salvos Launched in 256-K Battle," Electronics, June 30, 1983, p. 107.

Beresford, Roderic and Robert J. Kozuma, "Ferment and Recovery Stir U.S. Industry to Fight Japan," Electronics, June 30, 1983, p. 129.

Bishop, Jerry E., "Can a Supercomputer be Built? Team at IBM Grows More Confident," The Wall Street Journal, February 27, 1981.

Blodgett, Albert J., "Microelectronics Packaging," Scientific America, July 1983, p. 86.

"Britian Rises to Japan's Computer Challenge," Science, May 20, 1983, p. 799.

Bylinsky, Gene, "Japan's Ominous Chip Victory," Fortune, December 16, 1981, p. 52.

Bylinsky, Gene, "The Next Battle in Memory Chips," Fortune, May 16, 1983, p. 152.

"Chip Wars: The Japanese Threat," Business Week, May 23, 1983, p. 80.

Computers: 1971-1981, Center for Technological and Interdisciplinary Forecasting, Tel Aviv University, Ranat Aviv, 1971.

"Computers and Their Role in Energy Research--Current Status and Future Needs," Subcommittee on Energy Development and Applications and

Subcommittee on Energy Research and Production, U.S. Congress, House of Representatives, Washington, D.C., June 14, 1983.

"Data-Driven Automation: The Education Problem," Carol Truxel, ed., IEEE Spectrum, May 1983, p. 93.

Dohler, Gottfried, "Solid-state Superlattices," Scientific America, November 1983, p. 144.

Feigenbaum, Edward A. and Pamela McCorduck, The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World, Addison-Wesley Publishing Co., Reading, MA, 1983.

Ferguson, Charles H., "The Microelectronics Industry In Distress," Technology Review, August/September 1983, p. 24.

Gill, Michael D., Technology Venturing: The Japanese Fifth Generation Computer Program, unpublished Professional Report, The University of Texas at Austin, August 1983.

Goldhor, Richard S. and Robert T. Lund, "University to Industry Advanced Technology Transfer: A Case Study," Research Policy, June 1983, p. 121.

"Government Plan to Make Fast Computer With Business is Getting on Track," The Japan Economic Journal, July 20, 1982, p. 17.

ICOT Journal, Institute for New Generation Computer Technology, September 1983.

"Industry U.S. and Foreign," Science, October 21, 1983, p. 316.

Ingrassia, Lawrence, "Control Data and Cray Start Feuding Again," The Wall Street Journal, May 25, 1982, p. 31.

Iversen, Wesley R., "Supercomputers Find New Jobs," The Economist, July 28, 1982, p. 75.

"Japan Goes for the Gusto," Datamation, August 1981, p. 40.

"Japan Plans New Generation of Computers," New Scientist, July 16, 1981, p. 131.

Japanese Economic Journal, various issues, June 1980-October 1983.

Johnston, Bob, "Japan Builds a Pillar of the Fifth Generation," New Scientist, July 21, 1983.

Kolcum, Edward, "U.S. Plans Supercomputer for Defense," Aviation Week and Space Technology, April 18, 1983, p. 77.

Lehner, Urban C., "Japan Starting 10-Year Effort to Create Exotic Computer," The Wall Street Journal, September 23, 1982, p. 23.

Lester, Richard K., "Is the Nuclear Industry Worth Saving?" Technology Review, October 4, 1982, p. 39.

Levine, Ronald D., "Supercomputers," Scientific America, January 1982, p. 118.

Lohr, Steve, "Japan's New Test in Chips," New York Times, June 5, 1983, p. F1.

Malik, Rex, "Japan's Fifth Generation Computer Project," Futures, June 1983, p. 205.

Manuel, Tom, "Advanced Parallel Architectures Get Attention as a Way to Faster Computing," Electronics, June 16, 1983, p. 105.

March, Peter, "The Race for the Thinking Machine," New Scientist, July 8, 1982, p. 85.

"Microchip Battle: The 64K Battle," The Economist, May 15, 1982, p. 94.

Mowery, David C., "Economic Theory and Government Technology Policy," Policy Sciences, September 1983, p. 27.

Nakamura, Hideichiro, "Is Japan's Economic Strength Fortuitous?" Japanese Economic Studies, Spring 1983, p. 48.

"Now Japan is in the Running on Chip-making Machines," The Economist, January 15, 1983, p. 85.

"Outline of Research and Development Plans for Fifth Generation Computer Systems," Institute for New Generation Computer Technology, April 1983.

Pollack, Andrew, "Japan's Big Lead in Memory Chips," New York Times, February 28, 1982, p. F1.

Posa, John G., "What to Expect Next, A Special Report," Electronics, May 22, 1980, p. 119.

Preliminary Report on Study and Research on Fifth Generation Computers, 1979-1980, Japan Information Processing Development Center, 1981.

Reich, Robert B., The Next American Frontier, New York: Times Books, 1983.

Rostow, W. W., How it All Began: Origins of a Modern Economy, New York; McGraw Hill, 1975.

Schefter, Jim., "Fifth Generation Computers," Popular Science, April 1983, p. 79

Selizman, Mark, "The Fifth Generation," PC World, October 1983, p. 282.

"Silicon Valley Has a Big Chip About Japan," The Economist, March 20, 1982, p. 69.

"Slowly, Slowly on Chips," The Economist, March 27, 1982, p. 92.

Smith, Kevin, "U.K. Pursues Fifth Generation Computer," Electronics, May 31, 1983, p. 101.

"Straining to Lift Consumption After a Slow Start," Electronics, January 13, 1982, p. 121.

Summer, Larry W., "VHSIC: A Status Report," IEEE Spectrum, December 1982, p. 34.

"Supercomputing Research and Development, The LLNL Program and Proposals," unpublished manuscript, Lawrence Livermore National Laboratory, March 23, 1983.

"Technology Update: Information Technology," Electronics, October 6, 1983, p. 123.

"Telecommunications: The Global Battle," Business Week, October 24, 1983, p. 126.

Texas 2000 Commission, "Scenarios of Performance of the Energy Industry and the Information Technology Industry in Texas," unpublished manuscript prepared by The University of Texas at Austin Project Team.

Texas 2000 Commission, "Methodology Used in Forecasting Performance in the Information Technology and Energy Industries," unpublished manuscript prepared by The University of Texas at Austin Project Team.

Texas 2000 Commission, "Comparison of Historical Performance, Texas 100 Corporations vs. Standard and Poor's 500 Corporations," unpublished manuscript prepared by The University of Texas at Austin Project Team.

Texas 2000 Commission, "Historical Review and Forecasts of Performance in the National Energy Industry 1969-1990," unpublished manuscript prepared by The University of Texas at Austin Project Team.

Texas 2000 Commission, "Historical Review and Forecasts of Performance in the National Informational Technology Industry 1969-1990," unpublished manuscript prepared by The University of Texas at Austin Project Team.

"The Battle of the Supercomputers: Japan's All-out Challenge to the U.S.," Business Week, October 17, 1983, p. 156.

"The Datamation 100," Datamation, May 30, 1983, p. 97.

"The Race to Build a Supercomputer," Newsweek, July 4, 1983, p. 58.

Third Quarter Report, Intel Corporation, September 30, 1983.

"Two Nations Race to Build Super-speed Type Computers," The Japan Economic Journal, April 6, 1982, p. 9.

Uttal, Bro, "Here Comes Computer Inc.," Fortune, October 4, 1982, p. 82.

Verity, John W., "Endowing Computers With Expertise," Venture, November 1982, p. 48.

"Who Will Build the Next Supercomputer?" Science, January 1981, p. 268.

"Will Japan Leapfrog America on Superfast Computers?" The Economist, March 6, 1983, p. 95.

"Xerox Scientist Joins DOD Supercomputer Program," Science, June 24, 1983, p. 1359.

1982 Annual Report, Cray Research.

1982 10-K Report to the Securities and Exchange Commission, Cray Research.

Mr. Fuqua. Thank you very much, Dr. Kozmetsky, for a very thorough and, I think, thought-provoking essay on—or statement on supercomputers and where we're going.

Is it your opinion that the level of supercomputer research in our private industries, universities, and Government facilities, are adequate to provide for the future of supercomputing needs of the country? And—in other words, it is also desirable for the Federal Government to—to support all of these groups?

Dr. Kozmetsky. On your two-part question, for the first part, in my opinion, there are not sufficient supercomputers in our universities for them to discharge the function which the citizens of our country always wanted them to, which is to educate the people for a meaningful future. We're way behind now. So we have to make up lost time. And if it does take Federal support, then I think we should support it. I think we should certainly do it on a minimum of regions, and I think, in the name of equality, there should be a minimum of 20 established across this country as quickly as we can.

As to the second part, Mr. Chairman, I think it is absolutely imperative that we get our economy moving again, and the supercomputer is extremely important because it will be the driver that I think around these various areas will give birth to both small, medium-size—and as well as in the larger companies—entrepreneurship, just as we've seen at IBM, CDC, and some of the other companies.

Mr. Fuqua. Well, considering the economics—and you have some charts reflecting that of developing supercomputers, but I'm really speaking of supercomputer centers and the current availability of trained manpower—do you think that the number of federally initiated or supported supercomputer centers would—should be limited?

Dr. Kozmetsky. No, sir.

Mr. Fuqua. Well, what about the—do you see that we are adequate to meet the geographic and discipline and the training and the quality of service and other access concerns associated with these centers?

Dr. Kozmetsky. I think—first, let me give you my basic assumption. I don't think you can get everything moving at once, as you have indicated by your question. Second, if you don't have the ability to use the supercomputer, if you force the use of supercomputers to go into a priority or to get into a queue, you will break down its use. People will resort back to smaller computers, use what they now have, and we will fall further behind in many of the things that need to be done.

Just to give you those which I'm extremely concerned with, knowing, Mr. Chairman, your concern for energy, knowing my great State of Texas also is concerned about its wasting resources, I can't help but think that biomass is one of our answers in west Texas. We have no computers to simulate scaling up research so that we can form the new chemical industry of the United States for the 1990's or the 21st century.

I can tell you, sir, that if we are going to assist in medical care diagnosis and to be able to cut health costs, we desperately need supercomputers. And you can't queue each of these, they are all important.

Mr. **Fuqua**. Well, what's the—in accomplishing those objectives that you just outlined, what is the role of the Federal Government in that, and what is the appropriate role that the Federal Government should assume in something like that, if there is any? I assume there is some.

Dr. **Kozmetsky**. Yes; there is. And I think it should go back to the role where the Federal Government has always shown the leadership in our country for R&D. Most people have written on this. They've always stated that they pointed with great pride to NASA and to DOD. And I think there is a nonspace, nondefense aspect to it, in the sense that we are in a fierce competitive struggle. And that would be for us to put these computers out as quickly as we could and have them be funded by the Federal Government.

Mr. **Fuqua**. Well, what—what legislative initiatives would you think would be the most appropriate for the Government to pursue?

Dr. **Kozmetsky**. Well, I see three major things: first, the computers, which I've overemphasized; second, I'd like to see a delivery system, and obviously I would like to see it satellited to tie the higher educational institutions of this Nation together on that; third, establish one that would exchange software.

Mr. **Fuqua**. Would that be legislative initiatives that we've——

Dr. **Kozmetsky**. Yes. I think that could be done, and perhaps in partnerships with the States, because the States are also greatly concerned. And many—and there have been over 150 active programs taken by many States in our country to develop the high-tech industry, and I think this would be the role of the Federal Government in showing leadership, and certainly I would encourage matching funds between the two governments.

Mr. **Fuqua**. And that would be to increase the accessibility for our universities, and probably industry researchers to—by having the Federal Government stimulate that accessibility?

Dr. **Kozmetsky**. Yes. And I think that also entails the appropriate other agencies to make sure that—like the National Science Foundation encourages the research and development of various advanced applications, including the software, providing for scholarly exchanges during holidays and during the summers. So that we truly start building a community of scholars that are working closely with industry and Government in what I call technology venturing.

I think that will give rise and spinoff to a rebirth in this Nation that will make the days when I grieved so much when Sputnik went up, and I was on the team that mothballed the first two satellites in the United States 2 years previous to that. That the whole rebirth, the charge on it, it does not only include simply university people, sir. I would like to hasten—it provides for a replacement of blue collars with what I call the white smock industry, the maintenance people, the people that are going to be doing the detailing, the getting out, keeping track of things. It would be a tremendous revival.

Mr. **Fuqua**. Thank you.

Mr. Winn.

Mr. WINN. Thank you, Mr. Chairman.

Doctor, to what extent have you discussed the supercomputer design and operation with such computer centers, as the National Center for Atmospheric Research, magnetic fusion computer network, the University of Minnesota, Colorado State University, and with the industrial research consortia, such as the ETA Systems, MCC, and SRC? Their input would be, I think, very interesting and very helpful.

Dr. KOZMETSKY. I met with the chairman of coordinating all of the Government laboratory work. I have visited a half dozen of the Government laboratories personally. I have reviewed all of their annual reports. I certainly must commend them for their tremendous advances in scientific computation.

Sir, I'd like to humbly point out that we need to move from science to technology, which has always been a very valuable resource. What main emphasis I'm putting on dispersing the supercomputers as quickly as possible is to take advantage of American technology. And I'd like to hasten the process by putting things to work rather than printing, meeting——

Mr. WINN. Would you characterize your comments than as their chief concerns, too?

Dr. KOZMETSKY. It is their concern. But I think quite appropriately that they have priorities. Their priorities are the major things that they have been funded and discharged to do, both classified and unclassified, which gave birth and rise to the development and building of the first supercomputers in this country.

For example, sir, I'm very much involved, and have been for 6 years, in putting an extremely giant data base to answer such simple questions of what are the dying industries in the United States. I have a data base that contains 6,800 companies statistically comparable from 1950 to 1982, including all mergers and acquisitions put into it. I must do with some of the techniques, we evolved 5 million multiple correlations simultaneously. Sir, I can't get on any supercomputer in the United States to run that type of a problem.

Mr. WINN. Let me follow up on my first question, if I may, about concerns of these various people. Let's—let's really dream a little bit here and say that each of those groups that I mentioned had practically unlimited resources. Now what would their concerns be? Would they still be parochial, or do you think that they might want to go in the direction that you discussed or some other direction?

Dr. KOZMETSKY. I can tell you the contacts that I have made. Parochialism is almost disappearing from our major universities. We have seen a new generation of faculty. We have just an exciting new generation of students, both at the graduate and undergraduate level that want to get about the business and be preeminent. Pride is important.

Also, I have never seen an environment as we have today where we are eager to make sure that we do learn not only carrying science for the understanding of phenomenon, but there is more and more an awakening to carry the understanding for use.

Mr. WINN. The importance that you describe—ascribe to IBM compatibility appears either to forecast significant design changes by American supercomputer manufacturers or to anticipate their demise; am I reading you incorrectly?

Dr. KOZMETSKY. No, sir, I just simply said that the market, I think, will always be IBM compatible, and that includes IBM, sir. And I hope that—and I do think IBM will continue to make the United States a preeminent computer industry nation.

What I am saying is that, to date, other supercomputer manufacturers have not been IBM compatible. Just 2 weeks ago, both Hitachi and Fujitsu have entered into an agreement to license IBM software to be compatible. I have worked with Fujitsu through my boardship on Amdahl since 1970, and I can assure you that they're extremely aware what IBM compatible means.

They currently have taken their vector computer—or as we call it, the supercomputer—they have licensed it to ICL in England, and they have licensed it with Siemens in Germany. But both licenses at the current moment is for non-IBM compatible. They have yet to decide on what they're going to do with IBM compatible, either sell it themselves or let these other companies sell them in other parts of the world.

So the competition, I think, by 1984 is going to be visible and written up in the media.

Mr. WINN. You state that the two Japanese firms can deliver supercomputers to order today, and that one of these appears to have one-third to one-half higher capacity than the present American machines. What reaction from the American competitors do you hear?

Dr. KOZMETSKY. First, disbelief. They don't think that they have it, that they are not ready to deliver. Second, both the industry and the universities don't believe they have the software. I can assure you in my discussions with the executives of those companies that they do have it. I've already indicated that we have run one major large-scale manufacture—nondefense, nongovernment application on their computers.

Mr. WINN. So you figure that the status——

Dr. KOZMETSKY. I said they're ready to go. I've never found——

Mr. WINN [continuing]. Of the Japanese software is operable, is ready to go.

Dr. KOZMETSKY. Yeah, right. And also they have the IBM software and the compatibility between data bases.

Mr. WINN. Well, is the Japanese software, do you consider that IBM compatible?

Dr. KOZMETSKY. No; their operating software works on their machines. The IBM will work on their machines. Their machines are ready to deliver. I've been associated long enough with the Japanese to know they have never missed a delivery schedule that they've ever promised. They're ready to deliver.

Mr. WINN. Maybe we ought to send large delegations with cameras over to Japan to bring back some of their software.

Dr. KOZMETSKY. Well, let me explain myself. I don't want their software.

Mr. WINN. You want us to develop our own.

Dr. KOZMETSKY. I could put our software on their machines.

Mr. WINN. And get with it with Government help.

Dr. **Kozmetsky**. Yeah. And I can assure you, sir, some American companies don't need cameras. They have already entered into agreements a number of years ago, and already have licensed the Japanese advanced technology, and it is being manufactured already in California and sold worldwide. So at least one company has responded and is competing.

Mr. **Winn**. A last question, Mr. Chairman. I wonder if Dr. Kozmetsky could tell us what orders have been placed for near-term delivery from Japan? Do you know what those orders might be?

Dr. **Kozmetsky**. Certainly in their own country, they have started already to deliver to their universities. Which industries, I don't know. But I've been assured that they can deliver them quickly, and I wouldn't be surprised if they didn't continue, as they always have done. Fujitsu, in particular, has outsold IBM in Japan, and that's been true for the last 3, 4 years. I would not be surprised if their first deliveries, as they've always said, would first be in Japan. They've already gotten ready to enter the European market, and I wouldn't be surprised if they're still in the planning stage of trying to get into the American market.

Mr. **Winn**. Thank you very much.

Thank you, Mr. Chairman.

Mr. **Fuqua**. Mr. MacKay.

Mr. **MacKay**. Dr. Kozmetsky, that's—this is fascinating testimony. And I'd like to come back at the risk of some duplication and go into the question again that Chairman Fuqua asked.

What—do you see the role of Government in this as basically helping create a market to guarantee that the maximum number of American concerns will go for this market and, at the same time, as guaranteeing that we will have a trained source of talent——

Dr. **Kozmetsky**. Yes.

Mr. **MacKay**. Which zeros in on the university, a Government-university kind of coalition.

Dr. **Kozmetsky**. Let's see if I can't make that plainer and clearer.

If I bought a supercomputer and I could tie it in with my data processing and have my data bases and be ready to run on a supercomputer, I would then buy a supercomputer. So when I look at, say, the market for the 380 generation of the IBM machines, which is in the 6,000 to 10,000 range, that simply opens up a market for a supercomputer at our major companies.

If I were to go back to a company which won't be with us in 1984, January, the old AT&T, I imagine they bought someplace like 300 or more of these superlarge computers. And they would then be able to cut back on some of their data processing and put a supercomputer in. So that market could evolve extremely rapidly.

Now, what's lacking are people that know how to program it. And that's been our role in the universities. And we can't wait for the normal period of taking them through. So the more we could put in—and I can't—I could assure you how eager most universities are for other than the applications they've had. We all feel like we were left out, both hard science and the soft science people—that would then allow us quickly to get people that could be hired put in.

Also, it would encourage what I call technology venturing and something I know this committee has encouraged, getting the industry and the universities to work closer together, select things, set priorities. And that—and I think, in a short period, off we go. It would be almost as dramatic as going to the moon, and this time it would be better when it came back because we already helped planet Earth.

Mr. MacKay. That's—that's—so what we really need then is a commitment to a Federal—a commitment to a Federal role that says we will provide a market for supercomputers and it will be a stable, long-term, not a one-shot——

Dr. Kozmetsky. Well, I'm glad you brought that up again, because providing the market—you've stressed it too hard. We've got to develop the people that can be hired because of the market outside the universities. And that has to come first.

Mr. MacKay. Thank you.

Mr. Fuqua. Mr. Reid.

Mr. Reid. Thank you, Mr. Chairman.

I certainly applaud you for holding these hearings. This is something that we heard in one of the first hearings we held about supercomputers, and at that time it was only a word, and I've learned a little more about them since then.

Doctor, the question I have, though—you know, we watched your slide presentation. It was most interesting, and the work you did there is significant. But why, if you can project sales as you have, why have—why hasn't IBM or one of the other major companies said, OK, there's a big market for this, I'll go ahead and do it, we'll go ahead and do it?

Dr. Kozmetsky. Let's see if I can answer you succinctly, as I did probably separate it in the testimony.

First, if I were IBM, I'm positive that they started in 1981 when they announced that they were going to start building a computer out of gallium arsenide.

Mr. Reid. Going to start what?

Dr. Kozmetsky. Building a large computer using gallium arsenide as the basic component, which immediately takes you into a supercomputer.

So, first, I'm positive that IBM does have under development at some stage—and I don't know what stage it is—a supercomputer.

Second, when I take a look at all of the things that they're doing with personal computers, some of the things that they're doing to enhance the use of the other large-scale computer, their 380 series, and setting up the networks and everything else, when you take the percentage of it, it is a small percentage. As I stated as an economist, if I had a choice with scarce people where to put them, I'd put them where I'd get the highest marginal yield.

Mr. Reid. OK, I missed——

Dr. Kozmetsky. As a businessman, there's more gross margin on the other products than on a supercomputer. When that market builds, I'm positive, IBM will be right there, just as it did with the personal computer. It's a question of timing. And these next 5 years are very critical for us to get the infrastructure to use supercomputers and to compete.

Mr. Reid. OK. Go over this again with me. What do you propose we as a Congress—what should we do to make this 5 years something that's meaningful to us after the 5-year period's gone?

Dr. Kozmetsky. Encourage, in any way you can, the dispersal of supercomputers. I stuck my neck out for at least 20 across this country in institutions of higher learning with the appropriate infrastructure means. If we have to satellite it, let's satellite it. Let's network it so that we can get the maximum amount of our colleges working, getting used to it, advancing it, training people, because we are going to be using it. As I said earlier, it's a key resource——

Mr. Reid. In effect, what you're——

Dr. Kozmetsky. In new emerging industries that we all talk about and read in Megatrends and in the newspapers.

Mr. Reid. In effect, then, what you're saying is a small—relatively small investment by the Government today would pay dividends to all sectors of the economy in the future?

Dr. Kozmetsky. Right.

Mr. Reid. One last question.

A subcommittee of this full committee held hearings on computers in universities, and one of the things we heard there by all the people that testified is that, even though we have a lot of computers—and this, I am sure, would apply also to some of the testimony you've given—the software for the computers is just so bad, it's not interchangeable, there's nothing to put in them. Now wouldn't this be magnified significantly by the subject about which you're speaking?

Dr. Kozmetsky. No; I think I'd overcome that obstacle. Remember, I also said if we got 20 out, I'd also set up one more which would have all the software in it. So if I am at Austin and I want to go to Tallahassee and get software, I could zoom it out of a computer and put it into mine.

Mr. Reid. OK.

Dr. Kozmetsky. And I wouldn't have to reinvent it.

Mr. Reid. OK.

Thank you, Mr. Chairman.

Mr. Fuqua. Mr. MacKay.

Mr. MacKay. One technique that has been used at the State level with considerable success is, in effect, a kind of a challenge grant proposal, where I know in State legislative bodies, Florida specifically, they've put up 60 percent of the money for endowed shares, and they've told the universities, you match this with the other 40 percent. And it has had a real stimulative effect.

What would you think of a kind of a matched which, in effect, would be some Federal dollars, but also stimulating universities to get out and do a little self-help themselves?

Dr. Kozmetsky. Well, you couldn't have asked a more biased person, except I like about 50/50 better.

Mr. MacKay. Fifty/fifty is all right. That would be an interesting thing.

Mr. Fuqua. Thank you very much, Dr. Kozmetsky.

Dr. Kozmetsky. Thank you, Mr. Chairman. Thank you members.

Mr. Fuqua. We appreciate very much you being here.

At this time, I'd like to recognize—and I would just state that there'll be a vote momentarily—but I'd like to recognize our distinguished colleague from Indiana, John Myers, who has been a very valued Member, and one that this committee has worked with on his role on the Appropriations Committee on handling many of the various things we're talking about today. And we're very pleased to have our colleague, John Myers, present our next witness.

STATEMENT OF HON. JOHN T. MYERS, A REPRESENTATIVE IN CONGRESS FROM THE STATE OF INDIANA

Mr. MYERS. Thank you, Mr. Chairman, Mr. Winn, and members of the committee.

I first thank you very much for having this hearing today. I think it is extremely valuable that we have a concern about this. And the next witness, who is a new constituent of mine, has an extremely brilliant career in many fields, and I know that the testimony he's about to give will be very valuable to members of this committee.

Steve was born in Germany, and came into this country in 1948 as a 16-year-old immigrant. He has advanced very, very fast academically. He's a doctor now, and he's the kind of doctor that can help you, he's a medical doctor, having been most recently the dean of the Indiana University School of Medicine at Indianapolis, and this last summer became president of Purdue University, thereby becoming my constituent.

For interests of this committee, Purdue University has produced more astronauts than any other one university in our country.

He has already started out on a very brilliant career and has set very high goals for Purdue University, a very fine institution in Indiana and the world. So I know the testimony he's about to give will be very valuable, Dr. Steve Beering, president of Purdue.

Do you want to hear from him now?

Mr. FUQUA. We will go ahead and proceed until the bells ring.

Mr. MYERS. We are in conference with the Senate——

Mr. FUQUA. Yes.

Mr. MYERS. On a supplemental appropriations. I may stay, Steve, and listen.

Mr. FUQUA. Dr. Beering, we are very pleased to have you here. And you are also representing the American Association of Universities and the National Association of State Universities and Land-Grant Colleges. We are very pleased to have you here today.

STATEMENT OF DR. STEVEN BEERING, PRESIDENT, PURDUE UNIVERSITY, REPRESENTING THE ASSOCIATION OF AMERICAN UNIVERSITIES AND THE NATIONAL ASSOCIATION OF STATE UNIVERSITIES AND LAND-GRANT COLLEGES

Dr. BEERING. Thank you very much, Mr. Chairman, Congressman Myers, members of the committee. It is indeed a privilege to address you this afternoon on behalf of the 192 member universities of the AAU and NASULGC, and to follow the very articulate and exemplary testimony of the previous speaker.

Let me try to cast our testimony in a bit of a historical context. In the early 1800's, this country found itself at a significant crossroads in that, in order to advance further, we needed to have engineering and science, and we needed to have a vehicle to introduce the disciplines of engineering and science and agriculture into popular use.

So it occurred to our third President, Thomas Jefferson, to found a university for that purpose and, in 1819, when he did so, and founded the University of Virginia, he wrote in his preamble that this university would be different, that it should, "teach knowledge useful at this day." The subsequent history has, of course, demonstrated that he was right, and that this work in 1862, the Morril Act, generalized that philosophy and established the land grant university system throughout the country.

At that time, the crossroads were science and engineering. We are again at such a crossroads in America, moving from industrial to information society, from a hands-on to a heads-up technology situation. At this very time, we're being challenged by Germany, France, England, and Japan, in the development and application of the means that would assure our competitive edge in the international economy.

Nowhere is this problem more evident and poignant than in the supercomputer competition, as Dr. Kozmetsky so excellently described.

What are the components of the problem? I'll describe them in terms of six lacks—lack of access, lack of availability, lack of funds, lack of trained computer scientists, lack of software, and lack of communication linkages.

At the same time, there are significant opportunities. Again, the previous speaker alluded to many of them. These opportunities are not any longer confined to administrative number crunching, but they are evident in science and engineering, technology, agriculture, space exploration, medical research, and the military.

There are three major areas of opportunity. The first deals with simulation, such as you've already studied in the Department of Defense applications, the Department of Energy concerns, the petroleum industry such as seismic exploration, and the communications business and weather forecasting—but importantly also in chemical interactions and molecular and genetic analyses, which brings us right into one of the major medical problems, namely cancer research.

A second opportunity is modeling, the design of robots and equipment, development of new drugs, applications in the airline industry and in aerospace, complex real time operations such as data processing, the development of signal and image analyses, and again, a very important one in terms of medical application, namely the diagnosis of abnormal heart rhythms.

Having told you about these two applications, I've described the two major morbidity and mortality problems in this country; namely, cancer and heart disease.

The supercomputer then benefits not only in administrative computing, but also in research, in education and development and application, and, most importantly, as a Nation in maintaining its economic edge internationally.

Let me describe briefly the Purdue experience. Because of the extreme cost and some of the factors described by Dr. Kozmetsky, thus far, only three universities out of over 3,300 in this country have taken the expensive step to implement a supercomputer program.

We planned ours since 1979, and installed a Cyber 205 in January 1983. It became operational only 2½ months ago, on September 1, 1983. But already we have over 200 projects, mostly in engineering and science, and a number in agriculture. There's already a queue developing of people who want to get on the computer with new projects.

Importantly, we have decided to have all the programing done by graduate students under supervision of the faculty, because we believe that one of the great issues that we must resolve is the training of the computer scientist. We are, therefore, immediately benefiting not only research, but also instruction.

But perhaps the most exciting development is the translation of this new development into the economy. Therefore, we invited five major companies—in fact, we've just invited five more—to join us on campus in a program called CIDMAC, Computer Integrated Design Manufacturing Automation Center. These companies are currently involved with us in designing the factory of the future using the Cyber and robotics.

This is a project which is, in itself, exciting, and I would like to note that it is not only a technological advance, it is also a fine demonstration of interdisciplinary university work in that we are involving the schools of engineering, the school of science, the school of management, and then, of course, it is a fine example of university community collaboration which is actually crossing State lines, and that we're involving major corporations in this center.

A word about one other university that has a Cyber 205, that is Colorado State. They have been operational for a year, and they have elaborated a network involving seven universities, which demonstrates that universities can and do and will work together.

We would, therefore, on behalf of AAU and NASULGC, like to support the recommendations of the panel from which you heard last fall, chaired by Dr. Peter Lax and sponsored by DOE and NSF. This panel, as you may recall, recommended more access, more research, more training, and more development on a national, perhaps regional, basis.

In order to further elaborate on these points, I would like to point out to you some written testimony which is coming to your committee from the distinguished Hoosier and former colleague of yours, Dr. John Brademas, who is now the distinguished president of New York University.

In his testimony, which is designed to amplify mine, he makes five points, and I quote:

> The future development of computer technology in the United States requires cooperation among universities, industry and the government. Industry has a strong vested interest in vector processing, and needs incentives to broaden the prospects of new architectures. Venture capital interests have a two-year horizon, and because research and development time for supercomputers is likely to be longer term, venture capital is not likely to be a source of funds.
>
> Many imaginative designs for supercomputers are starting to emerge from university laboratories. We must find a way to support such efforts.

Universities also play an important role as a source of future researchers, computer architects, language and software designers, as well as users and developers of new application areas and methods.

There are several exciting projects now underway, some activity involving both industry and university participation. The success of these projects will eventually be tested in the marketplace, and will probably involve both existing as well as future companies.

And finally,

We believe that scientific

that is numerical—

machines, and thinking machines are not antithetical. A general-purpose computer well suited to both applications is possible. If such a machine is built, it will be attractive not only to researchers for the freedom it will give them to explore new ideas, but it will also be attractive to commercial developers as it will give them the widest possible world markets.

I would conclude then by saying that the development of large-scale computer capacity is clearly in the national interest. It is currently a scarce resource. In an investment in the future of America. It can be achieved by focused and concerted action, and collaboration of the public and private sectors, the government, the universities, business and industry is the way we would propose that we undertake it.

Thank you very much.

[The biographical sketch and prepared statement of Dr. Beering follow:]

OCS FILE COPY
SEP 1983

CURRICULUM VITAE

NAME: Steven Claus Beering

OFFICE ADDRESS:
President
Purdue University
West Lafayette, Indiana 47907
(317) 494-9708

HOME ADDRESS:
Westwood, 575 McCormick Road
West Lafayette, Indiana 47906
(317) 743-9933

DATE OF BIRTH: 20 August 1932; Berlin, Germany

WIFE: Jane (Pickering) Beering

CHILDREN:
Peter - June 22, 1960
David - April 7, 1962
John - May 28, 1969

HIGH SCHOOL: Taylor Allerdice High School; Pittsburgh, PA; 1950

COLLEGE: College of Liberal Arts; University of Pittsburgh, Pittsburgh, PA; Bachelor of Science, Summa cum Laude; 1954

MEDICAL SCHOOL: School of Medicine, University of Pittsburgh, Pittsburgh, PA; M.D., June, 1958

INTERNSHIP: Walter Reed General Hospital; Washington, D.C., 1958-59

RESIDENCY: Wilford Hall USAF Medical Center, San Antonio, TX; 1959-1962 (Internal Medicine)

MEDICAL LICENSURE:
Pennsylvania State Board License in Medicine and Surgery; August, 1959
Texas State License in Medicine and Surgery; 1967
Indiana State License in Medicine; 1969

SPECIALTY QUALIFICATIONS: Certified American Board of Internal Medicine; February, 1965

ACADEMIC AND ADMINISTRATIVE APPOINTMENTS:

- Professor of Medicine, Indiana University School of Medicine, 1969-
- Assistant Dean, Indiana University School of Medicine, 1969-70
- Associate Dean, Indiana University School of Medicine, 1970-74
- Director, Division of Postgraduate Medical Education, 1969-74
- Director, Indiana Statewide Medical Education System, 1970-83
- Coordinator, Indiana Regional Medical Program, 1972-76
- Dean, Indiana University School of Medicine, 1974-83
- Director, Indiana University Medical Center, 1974-83
- Chairman, Medical Education Board, State of Indiana, 1974-83
- Chairman, Commission on Medical Education, Indiana State Medical Association, 1973-83
- Secretary, Medical Distribution Loan Fund Board, State of Indiana, 1974-1981
- Secretary, The Attorney General of Indiana Public Health Trust Advisory Board, 1974-83
- Secretary, Medical Center Committee - I.U. Trustees, 1975-83
- Vice President, Medical and Nursing Distribution Loan Fund Board, 1981-83
- President, Purdue University, July 1, 1983-
- President, Purdue Research Foundation and Purdue University Foundations, 1983-

HOSPITAL APPOINTMENTS:

- Chief Resident in Medicine, Wilford Hall Medical Center, 1961-62
- Staff Internist, Endocrine-Metabolic Section, Wilford Hall Medical Center 1962-63
- Director of Education in Medicine, Wilford Hall Medical Center, 1963-67
- Consultant in Internal Medicine and Metabolic Disease - Santa Rosa Medical Center; San Antonio, 1966-69
- Chief, Internal Medicine Section, Wilford Hall Medical Center, 1967-69
- Attending Physician, Indiana University Hospitals, Wishard Memorial Hospital, Indianapolis V.A. Hospitals, 1969-
- Consultant Physician, Indianapolis V.A. Medical Center and St. Vincent Hospital, 1974-
- Chief of the Medical Staff, Indiana University Hospitals, 1974-83
- Medical Staff, Lafayette Home and St. Elizabeth Hospitals, 1983-

NATIONAL OFFICES:

- Chairman, Liaison Committee on Medical Education (International Accrediting Body for Medical Schools), 1976-81
- Vice President for USA Phi Rho Sigma Medical Society 1976-84
- Chairman, Section on Medical Schools of the American Medical Association, 1977-79; Councillor 1979-83
- Chairman, Council of Deans, Association of American Medical Colleges, 1980-81
- Chairman-elect, Association of American Medical Colleges, 1981-82
- Chairman, Association of American Medical Colleges, 1982-83
- Chairman, Council for Medical Affairs, (AMA, AAMC, AHA, ABMS, CMSS) 1981-83

HONORS AND AWARDS:

University of Pittsburgh Hall of Fame - 1954
Hooker Award in Pre-Medical Science - 1954
Owens Fellow in Medicine - 1955, 1956
Roche Award in Medicine - 1956
Heard Prize in Medicine - 1957
Wickerham Award in Medicine - 1957
Mosby Award in Medicine - 1958
Air Force Commendation Medal - 1964 (Medicine Teaching Program)
Distinguished Alumnus Award, Phi Rho Sigma Medical Society, 1969
The Surgeon General's Award for Distinction in Internal Medicine - 1969
Who's Who in America -
Who's Who in the World -
Sagamore of the Wabash, 1979-
Annual Beering Lecture - Indiana University School of Medicine
S. C. Beering Library, Indiana University Northwest Center for Medical Education
Steven C. Beering Professorship in Medicine, Indiana University School of Medicine, 1983-
Convocation Medal, American College of Cardiology, March 1983
University of Pittsburgh Distinguished Medical Alumnus Award (The Philip Hench Medal), 1983
Honorary Alumnus, Indiana University School of Medicine, 1983
Honorary Alumnus, Purdue University, 1983
National Academy of Sciences Institute of Medicine, 1983-

COMMITTEES AND CONSULTING ACTIVITIES:

Consultant to the Surgeon General (USAF) in Internal Medicine, 1965-69
Governor's Advisory Committee on Emergency Medical Services
Editorial Board, Journal of Indiana State Medical Association, 1970 to present
Governors' Committee on College Affairs, American College of Physicians, 1971-83
Consultant, Eli Lilly Research Laboratories, Indianapolis, 1971-83
Executive Council Association for Hospital Medical Education, 1975-78
Administrative Board, Council of Deans, AAMC, 1976-82
Counselor, Phi Rho Sigma Medical Society, 1976-84
Executive Council, AAMC, 1976-84
Indiana Statewide Health Coordinating Council, 1977-83
Association of American Medical Colleges Task Force on the Support (Financial) of Medical Education, 1977-80
Association of American Medical Colleges Task Force on Graduate Medical Education, 1977-80
Co-chairman, Indiana Cancer Crusade 1977-79
American College of Physicians, Committee on Educational Activities, 1979-82
Committee on Accreditation of Canadian Medical Schools, 1979-81
Health Advisory Council to the White River Park Commission, 1980-
American College of Physicians Committee on Health Policy, 1980-

COMMITTEES AND CONSULTING ACTIVITIES (cont.)

American Medical Association Council on Medical Education Task Force on "Future Directions for Medical Education," 1980-82
AAU Health Policy Committee, 1981-
Annenberg Center Academic Advisory Committee, 1983-
National Academy of Sciences/Institute of Medicine Committee on the Structure of the National Institute of Health, 1983-

BOARDS:

Board of Directors, Flanner House of Indianapolis, 1975-77
Board of Directors United Cerebral Palsy of Central Indiana, Inc., 1975-78
Regenstrief Foundation Board, 1975-
Krannert Institute of Cardiology Board, 1975-
Board of Directors, Greater Indianapolis Progress Committee, 1976-80
Board of Directors, American Cancer Society, 1976-83
University of Pittsburgh Board of Visitors, 1977-
Board of Directors, Indianapolis Convention and Visitors Association, 1979-
Board of Directors, Indianapolis Center for Advanced Research, 1979- ; Vice President 1982-
Board of Directors "Our House, Inc." - 1981-
Board of Directors, Eli Lilly and Company, 1983-
Indiana Economic Development Council, 1983-
Indiana Corporation for Science and Technology (also Executive Committee), 1983-

MILITARY SERVICE: USAF Medical Corps, May 1957 to June 1969

STATEMENT
FOR THE
SUPERCOMPUTER
HEARINGS

BEFORE THE

COMMITTEE ON SCIENCE AND TECHNOLOGY

U.S. HOUSE OF REPRESENTATIVES

BY

DR. STEVEN C. BEERING

PRESIDENT

PURDUE UNIVERSITY

ON BEHALF OF

ASSOCIATION OF AMERICAN UNIVERSITIES

NATIONAL ASSOCIATION OF STATE UNIVERSITIES
AND LAND-GRANT COLLEGES

NOVEMBER 15, 1983

SUPERCOMPUTER ISSUES FROM A UNIVERSITY PERSPECTIVE

TESTIMONY DELIVERED BY
STEVEN BEERING
PRESIDENT
PURDUE UNIVERSITY
BEFORE THE
HOUSE COMMITTEE ON SCIENCE AND TECHNOLOGY
TUESDAY, NOVEMBER 15, 1983

MY NAME IS STEVEN BEERING, AND I AM PRESIDENT OF PURDUE UNIVERSITY. I AM HONORED AND PLEASED TO HAVE THIS OPPORTUNITY TO APPEAR BEFORE YOU AND TO REPRESENT THE RESEARCH UNIVERSITIES AND COLLEGES WHICH COMPRISE THE MEMBERSHIP OF THE ASSOCIATION OF AMERICAN UNIVERSITIES AND THE NATIONAL ASSOCIATION OF STATE UNIVERSITIES AND LAND-GRANT COLLEGES. THESE ASSOCIATIONS REPRESENT THOSE UNIVERSITIES AND COLLEGES THAT CARRY OUT THE LARGER PART OF RESEARCH AND RELATED EDUCATION PROGRAMS SUPPORTED BY THE UNITED STATES GOVERNMENT. THIS AFTERNOON I SHALL PRESENT TO YOU OUR VIEWS REGARDING THE ISSUE OF UNIVERSITY ACCESS TO TODAY'S SUPERCOMPUTERS AND THE UNIVERSITY'S ROLE IN THE FUTURE DEVELOPMENT OF SUPERCOMPUTERS. ACCESS IS OF INCREASING IMPORTANCE TO RESEARCH UNIVERSITIES, AND I AM PLEASED THAT THIS COMMITTEE IS NOW PAYING ATTENTION TO THESE ISSUES.

CURRENTLY, ACCESS TO SUPERCOMPUTERS BY UNIVERSITY-BASED RESEARCHERS IS RESTRICTED DUE TO THREE FACTORS:

1) FEW RESEARCHERS HAVE THE TRAINING OR EXPERIENCE WITH WHICH TO USE THIS CLASS OF MACHINES EFFECTIVELY;

2) FEW INSTITUTIONS HAVE FUNDING AVAILABLE THROUGH GRANTS EITHER TO PURCHASE SUPERCOMPUTER SERVICES OR A SHARE OF SUCH RESOURCES; AND

3) FEW UNIVERSITIES HAVE SUPERCOMPUTERS ON THEIR CAMPUSES FOR USE IN RESEARCH AND TRAINING; IN ADDITION, FEW UNIVERSITIES POSSESS THE KIND OF COMMUNICATIONS FACILITIES THAT CAN PROVIDE ACCESS TO SUPERCOMPUTERS LOCATED ELSEWHERE.

SOME ADDITIONAL DETAIL WILL CLARIFY THESE POINTS.

A) NATIONAL SCIENCE FOUNDATION EMPHASIS ON AN "ENTREPRENEURIAL" APPROACH

DURING THE 1960S, THE NATIONAL SCIENCE FOUNDATION (NSF) CARRIED OUT A PROGRAM TO STIMULATE UNIVERSITIES TO EQUIP THEMSELVES WITH COMPUTERS TO SUPPORT TRAINING FOR AND USE OF THESE TOOLS IN EDUCATION AND RESEARCH ACTIVITY. NSF ASSUMED THAT, AFTER ITS PROVISION OF INITIAL CAPITAL STIMULUS, UNIVERSITIES WOULD FINANCE FURTHER GROWTH IN COMPUTER ACQUISITION THROUGH USER FEES. THEREFORE, AN ENTREPRENEURIAL, INSTITUTIONALLY-BASED APPROACH BECAME THE NATIONAL MODEL FOR ORGANIZING AND FINANCING ACADEMIC COMPUTING. ONE CONSEQUENCE OF THIS APPROACH WAS THAT NO UNIVERSITIES ACQUIRED SUPERCOMPUTERS UNTIL 1982, ALTHOUGH THERE WERE A FEW EFFORTS (SUCH AS THAT AT THE NATIONAL CENTER FOR ATMOSPHERIC RESEARCH) TO DEVELOP DISCIPLINE-BASED SUPERCOMPUTER FACILITIES FOR ACADEMIC USE.

THE NSF'S ENTREPRENEURIAL APPROACH, AIDED AND ABETTED BY RESTRICTIVE FEDERAL POLICIES FOR COST ALLOCATION, HAS NOT PROVED AN EFFECTIVE APPROACH FOR USE OF THE SUPERCOMPUTER BY ACADEMIC INSTITUTIONS AT THE LEADING EDGES OF SCIENTIFIC RESEARCH.

B) LIMITED UNIVERSITY EXPERIENCE

MOST EXISTING SUPERCOMPUTERS HAVE BEEN LOCATED IN INDUSTRIAL RESEARCH FACILITIES OR FEDERAL LABORATORIES. THUS THE POOL OF UNIVERSITY RESEARCHERS WHO COULD FORMULATE PROBLEMS FOR SUPERCOMPUTER ANALYSIS HAS BEEN LIMITED; SIMILARLY, THE NUMBER OF GRADUATE STUDENTS WHO COULD BE OR HAVE BEEN TRAINED TO USE SOFTWARE NEEDED FOR SUPERCOMPUTERS HAS BEEN LIMITED. MOREOVER, THIS LACK OF DIRECT ACCESS HAS ALSO RESTRICTED THE NUMBER OF GRANT REQUESTS WHICH HAVE INVOLVED USE OF SUPERCOMPUTERS.

C) STRUCTURE OF THE NATIONAL SCIENCE FOUNDATION

COMPUTING NEEDS SPAN A WIDE RANGE OF DISCIPLINES. THE NEED FOR THE SUPERCOMPUTER RANGES ACROSS THE ENTIRE ORGANIZATIONAL STRUCTURE OF NSF; YET THERE IS NO ORGANIZATIONAL UNIT WITHIN THE AGENCY WHICH IS RESPONSIBLE FOR THE IDENTIFICATION OF COMPUTING NEEDS IN SUPPORT OF RESEARCH GOALS. THE RECENT NSF REPORT, "A NATIONAL ENVIRONMENT FOR ACADEMIC RESEARCH," INDICATES THAT AN EFFECTIVE PROGRAM FOR ADDRESSING THE USE IN RESEARCH OF SUPERCOMPUTERS AND ASSOCIATED SUPPORT TECHNOLOGIES (SUCH AS NATIONAL COMMUNICATION NETWORKS) WILL REQUIRE A DIFFERENT APPROACH BY NSF. ONE SHOULD NOTE THAT NSF DID HAVE AN OFFICE OF COMPUTING ACTIVITIES DURING THE 1960S WHICH FOCUSED BOTH ON COMPUTING PROGRAMS AND ON RESEARCH IN COMPUTER SCIENCE.

D) SHIFTS IN COMPUTER CAPABILITY: MICROCOMPUTERS ON CAMPUS

A SIGNIFICANT FACTOR LEADING TO CHANGES IN THE USE OF SHARED CAMPUS RESOURCES FOR COMPUTING HAS BEEN THE DEVELOPMENT OF POWERFUL MICROCOMPUTERS. IN THE PAST, CAMPUS COMPUTING WAS TWO-TIERED IN NATURE: A CENTRAL CAMPUS FACILITY, USUALLY SUPPORTING TERMINAL ACCESS; AND INDIVIDUAL MICROCOMPUTERS FOR

USE IN DATA ACQUISITION AND EXPERIMENTAL CONTROL. CURRENTLY, CAMPUS COMPUTING INVOLVES TWO ADDITIONAL TIERS: THE SUPER-MINICOMPUTER, USUALLY SHARED BY TWO OR MORE DEPARTMENTS OR BY GROUPS OF RESEARCHERS; AND PERSONAL COMPUTERS. ACADEMIC RESEARCHERS HAVE TENDED, THROUGH USE OF THEIR GRANT FUNDS, TO SUPPORT ACQUISITION OF SUPER-MINICOMPUTERS. INSTITUTIONS AS A WHOLE, HOWEVER, HAVE EXPERIENCED GREAT DIFFICULTY IN GAINING ACCESS TO EXTERNAL GRANT FUNDS FOR SUPPORTING ACQUISITION AND SUPPORT OF CAMPUS-WIDE CENTRAL COMPUTERS.

E) LIMITED USE OF EXISTING TELECOMMUNICATION NETWORKS

THE SUPERCOMPUTER ISSUE IS LINKED TO THE NEED FOR EFFECTIVE COMMUNICATION NETWORKS TO LINK UNIVERSITIES AND SUPERCOMPUTER SITES. PARTICULARLY DURING THE TRAINING AND DEVELOPMENT PERIOD, UNIVERSITIES NEED DIRECT ACCESS TO SUPPORT STAFF AT SUPERCOMPUTER INSTALLATION SITES. IF COMMUNICATIONS BAND WIDTHS WERE SUFFICIENT, A LARGER CLASS OF ACTIVITIES COULD BE CARRIED OUT REMOTE FROM THE SITE. THE RESEARCH COMMUNITY HAS A GROWING NEED FOR RAPID DISSEMINATION AND EXCHANGE OF INFORMATION. SOFTWARE, DATA BASES, AND GRAPHICS ARE BECOMING A CRITICAL ELEMENT OF THE RESEARCH PROCESS. LACK OF CONVENIENT AND HIGH SPEED NETWORKS, THEREFORE, WILL LIMIT RESEARCH ACTIVITY. BECAUSE OF THEIR MASSIVE CAPACITIES FOR INPUT AND OUTPUT FLOWS, SUPERCOMPUTERS MAY PROVIDE A MAJOR RATIONALE FOR DEVELOPMENT OF SIGNIFICANT NATIONAL PROGRAMS TO ADDRESS THE NEED FOR EFFECTIVE COMMUNICATIONS NETWORKS. IT IS TRUE THAT A VARIETY OF NETWORKS ARE ALREADY IN USE WITHIN THE ACADEMIC COMMUNITY, SUCH AS ARPA, CSNET, EDUNET, AND THE LIKE. BUT THESE NETWORKS ARE STILL TOO LIMITED TO MEET THE NEED I AM ADDRESSING HERE.

AFTER LISTING ALL THESE FACTORS WHICH LIMIT UNIVERSITY ACCESS TO AND USE OF

SUPERCOMPUTERS, I SHOULD NOTE AGAIN THAT SEVERAL UNIVERSITIES HAVE TAKEN SIGNIFICANT FINANCIAL RISKS IN ORDER TO ACQUIRE SUPERCOMPUTERS. IN THE DEVELOPMENT OF ANY NATIONAL PROGRAM OR INITIATIVE TO ADDRESS THE PROBLEM OF ACCESS TO SUPERCOMPUTERS, IT IS ESSENTIAL THAT THESE INSTITUTIONS SHOULD BE SEEN AS CRITICAL NATIONAL RESOURCES AND THAT THEY SHOULD BE PROVIDED WITH ADEQUATE SUPPORT TO MAINTAIN SUCH RESOURCES FOR USE BY THE LARGER RESEARCH COMMUNITY. IN ANY PROGRAM FOR EXPANDING ACCESS TO SUPERCOMPUTERS, AN ATTEMPT SHOULD BE MADE TO BALANCE GROWTH IN ACQUISITION OF THESE INSTRUMENTS BY FEDERAL LABORATORIES WITH THE NEED TO LOCATE SOME FACILITIES IN UNIVERSITIES TO ADDRESS NEEDS FOR RESEARCH, GRADUATE TRAINING, AND THE DEVELOPMENT OF NEW ALGORITHMS AND SOFTWARE.

PURDUE UNIVERSITY IS ONE OF THOSE FEW INSTITUTIONS WHICH HAVE INVESTED IN SUPERCOMPUTING. IN THE YEARS 1979 TO 1981 THERE WERE A NUMBER OF MAJOR PLANNING ACTIVITIES AT PURDUE TO STUDY THE COMPUTING NEEDS OF THE UNIVERSITY IN THE AREAS OF RESEARCH AND INSTRUCTION. ONE OF THE NEEDS THAT WAS RECOGNIZED WAS THE NEED FOR SUPERCOMPUTER ACCESS FOR THE RESEARCH PROGRAM AND FOR RELATED INSTRUCTIONAL PROGRAMS AT PURDUE. DURING MOST OF THE PLANNING PHASE, IT APPEARED THAT THE ONLY SUCH ACCESS WOULD BE THROUGH THE USE OF SUPERCOMPUTER FACILITIES AT THE NATIONAL LABORATORIES MAINTAINED BY NSF AT NCAR AND BY THE DEPARTMENT OF ENERGY AT LOS ALAMOS AND AT OTHER SITES.

OUR PLANNING STUDIES CONVINCED US THAT THERE WAS A LARGE AND GROWING NEED FOR SUPERCOMPUTER ACCESS IN ALL OF THE DEPARTMENTS IN THE SCHOOLS OF SCIENCE AND ENGINEERING AND IN SOME OF THE DEPARTMENTS OF OTHER SCHOOLS SUCH AS AGRICULTURE, HUMANITIES AND SOCIAL SCIENCES, AND MANAGEMENT. IN 1980 TO 1981, IT WAS CLEAR THAT, FOR THE THEN-FORESEEABLE FUTURE, ONLY A SMALL FRACTION OF

THIS NEED COULD BE SATISFIED THROUGH THE NATIONAL LABORATORIES. THE DECISION WAS MADE TO PURCHASE A CYBER 205 SYSTEM FROM CDC. AN ORDER WAS PLACED IN NOVEMBER, 1981, AND THE SYSTEM WAS DELIVERED AND INSTALLED IN JANUARY, 1983.

OUR PLANNING PREDICTIONS THAT THERE WOULD BE CONSIDERABLE DEMAND AMONG PURDUE RESEARCHERS FOR ACCESS TO THE CYBER 205 HAS BEEN BORNE OUT. THE CYBER 205 WAS OPENED TO LIMITED PRODUCTION USE ABOUT ONE MONTH AFTER IT WAS INSTALLED, AND A NUMBER OF USERS WERE ALREADY PREPARED TO UTILIZE IT.

FULL-TIME USE ON THE 205 STARTED ON SEPTEMBER 1, 1983. BY THE TIME THIS IS BEING WRITTEN, AT THE END OF OCTOBER, 1983, THERE ARE ABOUT 150 RESEARCH ACCOUNTS ON THE CYBER 205. ALL OF THESE ACCOUNTS REPRESENT RESEARCH PROJECTS CARRIED ON BY FACULTY AND STAFF AT PURDUE. AS EXPECTED, MOST OF THESE ACCOUNTS (80) ARE IN THE SCHOOL OF ENGINEERING, AND MOST OF THE REST (55) ARE IN THE SCHOOL OF SCIENCE. THERE ARE NINE ACCOUNTS IN THE SCHOOL OF AGRICULTURE AND JUST A FEW IN OTHER AREAS OF THE UNIVERSITY. MOST OF THE ACTUAL PROGRAMMING IS BEING DONE BY GRADUATE STUDENTS, AND THERE WILL BE LITERALLY HUNDREDS OF GRADUATE STUDENTS AT PURDUE TRAINED IN THE APPLICATIONS OF VECTOR COMPUTING DURING THIS FIRST ACADEMIC YEAR IN WHICH THE 205 IS AVAILABLE ON A FULL-TIME BASIS. THE MAJOR CONTRIBUTION OF THE CYBER 205 IN THE AREA OF INSTRUCTION WILL BE IN CONNECTION WITH THE RESEARCH CARRIED ON BY THESE GRADUATE STUDENTS. HOWEVER, THE CYBER 205 IS ALSO AVAILABLE FOR STUDENT PROJECTS IN THOSE COURSES IN WHICH IT IS APPROPRIATE TO USE A VECTOR SUPERCOMPUTER. A FEW ADVANCED COURSES IN ENGINEERING AND IN COMPUTER SCIENCE ARE ALREADY USING THE CYBER 205.

WE STRONGLY RECOMMEND THAT SUPPORT FOR SUPERCOMPUTING IN UNIVERSITY COMPUTING CENTERS BE MADE PART OF THE NSF PROGRAM TO FACILITATE ACCESS TO SUPERCOMPUTERS FOR THE UNIVERSITY SCIENTIFIC AND ENGINEERING COMMUNITY. THERE

ARE NOW A NUMBER OF MAJOR UNIVERSITIES, INCLUDING PURDUE, THAT PROVIDE AN EXCELLENT COMPUTING ENVIRONMENT FOR BOTH RESEARCH AND INSTRUCTION. SUCH UNIVERSITIES SHOULD BE CONSIDERED AS APPROPRIATE SITES FOR NEW AND ENHANCED SUPERCOMPUTER INSTALLATIONS TO SUPPORT VERY LARGE-SCALE COMPUTING FOR UNIVERSITY SCIENTISTS AND ENGINEERS. SUCH UNIVERSITIES PROVIDE A TOTAL ENVIRONMENT DEDICATED TO RESEARCH AND INSTRUCTION, AN ENVIRONMENT THAT IS MUCH BETTER SUITED TO THE PURPOSES AND GOALS OF THIS NSF PROGRAM THAN, FOR EXAMPLE, THE VERY RESTRICTED ENVIRONMENT OF A WEAPONS LABORATORY. VECTOR SUPERCOMPUTERS IN THE UNIVERSITY ENVIRONMENT WILL ALSO HELP PRODUCE THE NEXT GENERATION OF RESEARCHERS AND TEACHERS WHO WILL BE TRAINED TO USE AND TO TEACH OTHERS TO USE MODERN LARGE-SCALE COMPUTING.

THE FUTURE DEVELOPMENT OF COMPUTER TECHNOLOGY IN THE UNITED STATES REQUIRES COOPERATION BETWEEN UNIVERSITIES, INDUSTRY AND THE FEDERAL GOVERNMENT. INDUSTRY HAS A STRONG VESTED INTEREST IN "VECTOR PROCESSING" AND NEEDS INCENTIVES TO WIDEN THE VIEW TOWARDS NEW ARCHITECTURES. VENTURE CAPITAL INTERESTS HAVE A TWO YEAR HORIZON, AND AS RESEARCH AND DEVELOPMENT TIME FOR SUPERCOMPUTERS IS LIKELY TO BE LONGER TERM, VENTURE CAPITAL IS NOT LIKELY TO BE A SOURCE OF FUNDS.

THERE ARE A NUMBER OF EXCITING PROJECTS NOW UNDERWAY, SOME ACTIVELY INVOLVING BOTH INDUSTRY AND UNIVERSITY PARTICIPATION. THE REAL SUCCESS OF THESE PROJECTS WILL EVENTUALLY BE TESTED IN THE MARKETPLACE AND WILL PROBABLY INVOLVE BOTH EXISTING AS WELL AS FUTURE COMPANIES.

WE BELIEVE THAT "SCIENTIFIC" (I.E. NUMERICAL) MACHINES AND "THINKING" MACHINES ARE NOT ANTITHETICAL. A GENERAL-PURPOSE COMPUTER WELL SUITED TO BOTH

APPLICATIONS IS POSSIBLE. IF SUCH A MACHINE IS BUILT, IT WILL BE ATTRACTIVE NOT ONLY TO RESEARCHERS FOR THE FREEDOM IT WILL GIVE THEM TO EXPLORE NEW IDEAS, BUT WILL ALSO BE VERY ATTRACTIVE TO COMMERCIAL DEVELOPERS, AS IT WILL GIVE THEM THE WIDEST POSSIBLE MARKET.

IN SUMMARY:

--CURRENT USE OF SUPERCOMPUTERS FOR RESEARCH BY UNIVERSITIES IS LIMITED.

--BARRIERS TO UNIVERSITY USE HAVE INCLUDED TRAINING, SOFTWARE AND ALGORITHM DEVELOPMENT, AND MACHINE CYCLES, AS WELL AS LIMITATIONS IN AVAILABLE CAPITAL RESOURCES.

--THE ENTREPRENEURIAL MODEL FOR FEDERAL FUNDING OF COMPUTING HAS PLAYED A SIGNIFICANT PART, IN RESTRICTING THE NUMBER OF SUPERCOMPUTERS IN UNIVERSITIES.

--THE NATIONAL SCIENCE FOUNDATION CURRENTLY LACKS AND SHOULD ADDRESS DEVELOPMENT OF AN EFFECTIVE ORGANIZATIONAL STRUCTURE FOR SUPPORT OF A PROGRAM FOR RESEARCH ACCESS TO SUPERCOMPUTERS.

--EXPANSION OF SUPERCOMPUTER RESOURCES IN FEDERAL LABORATORIES MUST BE COUPLED WITH A PROGRAM OF INSTALLATION AND REGIONAL SHARING OF SUPERCOMPUTERS AT UNIVERSITIES, INCLUDING AS A BASE THE SUPERCOMPUTERS WHICH HAVE ALREADY BEEN ACQUIRED BY UNIVERSITIES.

--DEVELOPMENT OF ANY NATIONAL SUPERCOMPUTER PROGRAM MUST ALSO INCLUDE CONSIDERATION OF THE RELATION OF SUPERCOMPUTERS IN INDUSTRIAL LABORATORIES AND UNIVERSITY RESEARCH NEEDS.

FINALLY, I WOULD SUPPORT THE RECOMMENDATIONS MADE THIS PAST DECEMBER BY THE PANEL ON LARGE SCALE COMPUTING IN SCIENCE AND ENGINEERING, SPONSORED BY NSF AND THE DEPARTMENT OF DEFENSE AND CHAIRED BY PETER D. LAX. THAT PANEL RECOMMENDED A LONG-TERM NATIONAL PROGRAM ON LARGE SCALE COMPUTING WITH PARTICIPATION BY APPROPRIATE FEDERAL AGENCIES, UNIVERSITIES, AND INDUSTRY, TO ADDRESS FOUR FUNDAMENTAL ISSUES:

1) INCREASED ACCESS FOR THE SCIENTIFIC AND ENGINEERING RESEARCH COMMUNITY THROUGH HIGH BAND WIDTH NETWORKS TO ADEQUATE AND REGULARLY UPDATED SUPERCOMPUTING FACILITIES AND EXPERIMENTAL COMPUTERS;

2) INCREASED RESEARCH IN COMPUTATIONAL MATHEMATICS, SOFTWARE, AND ALGORITHMS NECESSARY TO THE EFFECTIVE AND EFFICIENT USE OF SUPERCOMPUTER SYSTEMS;

3) TRAINING OF PERSONNEL IN SCIENTIFIC AND ENGINEERING COMPUTING; AND

4) RESEARCH AND DEVELOPMENT BASIC TO THE DESIGN AND IMPLEMENTATION OF NEW SUPERCOMPUTER SYSTEMS OF SUBSTANTIALLY INCREASED CAPABILITY AND CAPACITY BEYOND THAT LIKELY TO ARISE FROM COMMERCIAL SOURCES.

INCREASING ACCESS BY UNIVERSITIES TO SUPERCOMPUTERS IS IN THE NATIONAL INTEREST, CAN BE MET WITH A NATIONAL APPROACH, AND REQUIRES A NATIONAL PROGRAM.

Mr. **Fuqua**. Thank you very much, Dr. Beering. We appreciate your statement, and we will make the correspondence from our former colleague, John Brademas, part of the record following the questions and answers of Dr. Beering.

What problems do you see as far as in the supercomputers doing work for the Federal Government under certain procedures, particular the—the impact of the circular A-21? Do you see that it poses any barriers on cost recovery and resource sharing with—with universities, particularly in using supercomputers?

Dr. **Beering**. Well, I—I think we've worked out the mechanisms to do work for the Government, both on a contract and a grant basis, and I wouldn't see having a new gadget like this making it necessary to change our modus operandi. I think it would work just fine.

Mr. **Fuqua**. So you don't see that as any impediment?

Dr. **Beering**. No.

Mr. **Fuqua**. The—at the Livermore Lab in California, the Magnetic Fusion Computer Center is considered a successful model for supercomputing facililities. One of the components of this success is its software development and sharing that—is available for future—for fusion researchers. What kind of capabilities do universities have for developing and disseminating this kind of specialized hardware?

Dr. **Beering**. Well, I think we obviously aren't in the weapons business or the fusion research in most instances. Our major contribution would be the training of people and doing research, and perhaps in translating this work with business and industry.

Mr. **Fuqua**. Mr. Winn.

Mr. **Winn**. Thank you, Mr. Chairman.

I have one—I have several questions, Dr. Beering, that I'd like to submit for the record because we have a very important vote here very shortly on ERA. But I'm intrigued with the fact that Purdue, and Colorado State and the University of Minnesota are three universities that have bought the supercomputers without Federal help. Can you tell me how Purdue got theirs, how they got it started, and do you know how the other two schools did?

Dr. **Beering**. I don't know how the others did. I can tell you about Purdue. We happen to have a strong partnership with our State and with the other universities in the State. And we were able to get $5½ million from the State as seed money, and then the rest generated through our internal resources. It is a very expensive undertaking, as you know. The least of figures of such a thing, we are talking of $10 to $20 million just to get setup.

Mr. **Winn**. So it wasn't 100 percent funded by——

Dr. **Beering**. That's right.

Mr. **Winn** [continuing]. The State of Indiana.

You commented in your testimony of the importance of NSF in providing greater access to supercomputers for universities, and now we've just discussed that there was no Federal funding involved in the three schools that we've mentioned. These actions by universities have placed some pressure on the university system, I suppose, to make the programs work. What then is the proper role for the Federal Government——

Dr. **Beering**. Well, I would like to——

Mr. **Winn** [continuing]. NSF, DOE, and you referred——

Dr. **Beering**. I would like to second what Dr. Kozmetsky already alluded to, and that is that this is such a major national and even international problem that we ought to set aside any personal and parochial and individual views and make it a national effort, and make it a multifaceted national effort, and anyone who can play ought to be encouraged to do so.

As far as the universities are concerned, it would be best if we went through our traditional relationships that we have with NIH and NSF and the contract mechanisms with DOE and DOD, and so on, and if you were to fund those organizations and allow intensification of these relations.

Mr. **Winn**. You mentioned the 200 customers—I think that's what you said—or programs, I guess you call them.

Dr. **Beering**. Internal, yes.

Mr. **Winn**. Internal. Out of what? What—of course, I realize that it depends on the number and the time of the program. But what do you think your capacity or your——

Dr. **Beering**. Well, what we anticipate——

Mr. **Winn** [continuing]. Potential is.

Dr. **Beering**. We anticipate is rather grandiose, perhaps, but now possible. We have 52,000 students on 5 campuses, and this program is housed on the main campus. We want to connect, first of all, the entire State.

By the way, we want to do something about our library resources by going via computer instead of just buying more volumes and more shelves.

We want to take all of the students in science and engineering and technology and have them access this directly. And we already have some 5,000 students who are doing that. That's 5,000 students who are actually doing that, who are accessing the Cyber right now.

We would like to consider computers as a liberal education for every student, and feel that, before the decade is out, that every student, regardless of his major, is going to take at least one course in computers. And again we are gratified by the number of students in other disciplines, in the humanities and you name it, throughout the school who are interested in the survey courses that are now available.

So we envision that our capacity needs to build as a minimum on about 32,000 regular students, plus, of course the faculty.

Mr. **Winn**. What guidance would you give to other colleges and universities which would hope for supercomputer possession?

Dr. **Beering**. Well, we would hope that they would not be scared by it and that they would get into it in an incremental fashion as we have done. And that means, of course, that you first get into the computer game. You don't start off with a supercomputer. And there is a great use for the fifth generation computers as an ancillary device.

Mr. **Winn**. Thank you, Mr. Chairman.

I would like to submit a few other questions along the same line of questioning that I started.

Mr. **Fuqua**. Thank you, Mr. Winn.

Mr. MacKay.

Mr. **MacKay**. Mr. Chairman, I'll be very brief.

Dr. Beering, I think you have centered on one point that this committee should follow up, and that is that the NSF is not currently structured, because of its internal structure historically, so that the importance of a multidiscipline issue like this can be given priority.

I'd like you, if you would, to give us your suggestions, and possibly Dr. Kozmetsky also, as to what structure there should be. And I'd like to suggest to the chairman and staff director that we might continue this discussion with NSF.

It seems to me that us coming in now and putting in another institution to try to get money there when we've already got one might end up being counterproductive. And it might be better to zero in on the question of whether there are structural—how would you say it? It has become obsolete. The structure has become obsolete because of the new technology.

Mr. **Fuqua**. I might point out we do have Dr. Knapp who will be testifying tomorrow——

Mr. **MacKay**. OK. That's good.

Mr. **Fuqua** [continuing]. The Director of NSF, as well as Dr. Keyworth, tomorrow——

Mr. **MacKay**. Excellent. Then we can just discuss that directly.

Mr. **Fuqua**. And this can be explored with them.

Mr. **MacKay**. I'd like to thank both of you. This has been very stimulating. Thank you.

Mr. **Fuqua**. Thank you.

And thank you very much, Dr. Beering. We appreciate your being here today and your testimony.

Thank you, John Myers for being here.

The meeting stands adjourned until tomorrow morning.

[Whereupon, at 2:55 p.m., the committee was adjourned.]

SUPERCOMPUTERS

WEDNESDAY, NOVEMBER 16, 1983

House of Representatives,
Committee on Science and Technology,
Washington, D.C.

The committee met, pursuant to call, at 10:25 a.m., in room 2318, Rayburn House Office Building, Hon. Don Fuqua (chairman of the committee) presiding.

Present: Representatives Fuqua, MacKay, Winn, Lujan, Carney, McGrath, Skeen, Lowery, Boehlert, and Lewis.

Mr. **Fuqua**. The committee will be in order.

We begin today the second day of hearings on the future of supercomputers and relationship and role of the Federal Government in that. This morning we will begin with Dr. Edward Knapp, the Director of the National Science Foundation, and he will be followed by Dr. Edith Martin, the Deputy Under Secretary of Research and Advanced Technology with the Department of Energy. If Dr. Knapp and Dr. Martin would come forward.

Dr. Knapp, we are happy to have you back once again. We welcome—before we begin with Dr. Knapp, I would like to yield to Mr. Winn.

Mr. **Winn**. Thank you, Mr. Chairman.

Today's hearings on supercomputers is the fourth in nearly 6 months, and I predict that it will be equally helpful. Throughout this series we have heard from highly capable and knowledgeable people. This session is no exception.

Dr. DeVries on the President's Science Advisory Staff, Dr. Martin, Deputy Under Secretary for Research and Advanced Technology at DOD, and Dr. Knapp, the Director of the National Science Foundation, are administrators well aware of the important role which supercomputers presently fulfill and will exercise in the future.

[The prepared statement of Mr. Winn follows:]

THE HONORABLE LARRY WINN, JR.
HEARING ON SUPERCOMPUTERS
NOVEMBER 16, 1983

THANK YOU, MR. CHAIRMAN. TODAY'S HEARING ON SUPERCOMPUTERS IS THE FOURTH IN NEARLY SIX MONTHS AND I PREDICT THAT IT WILL BE EQUALLY HELPFUL. THROUGHOUT THIS SERIES WE HAVE HEARD FROM HIGHLY CAPABLE AND KNOWLEDGEABLE PEOPLE. THIS SESSION IS NO EXCEPTION. DR. RALPH DEVRIES, ON THE PRESIDENT'S SCIENCE ADVISORY STAFF; DR. MARTIN, DEPUTY UNDERSECRETARY FOR RESEARCH AND ADVANCED TECHNOLOGY AT DOD; AND DR. KNAPP, DIRECTOR OF THE NATIONAL SCIENCE FOUNDATION, ARE ADMINISTRATORS WELL AWARE OF THE IMPORTANT ROLES WHICH SUPERCOMPUTERS PRESENTLY FULFILL AND WILL EXERCISE IN THE FUTURE.

MR. CHAIRMAN, WE CONTINUE TO HEAR THE SAME MESSAGES FROM ALL OUR WITNESSES. THEY SUMMARIZE AS FOLLOWS: ONE, SUPERCOMPUTERS ARE NOT EQUALLY AVAILABLE TO ALL OUR SCIENCE AND ENGINEERING INVESTIGATORS AND TWO, THE UNITED STATES IS IN GRAVE DANGER OF LOSING ITS INTERNATIONAL POSITION OF SUPERCOMPUTER SUPERIORITY. I WILL NOT ENLARGE ON THESE STATEMENTS; WE ARE ALL FAMILIAR WITH THE CLAIMED RAMIFICATIONS AND CONSEQUENCES. I MUST CLAIM SOME DISAPPOINTMENT THAT WE HAVE HEARD SO LITTLE ABOUT HOW WE GOT TO WHERE WE ARE NOW. IT SEEMS TO ME THAT THIS COMPUTER CAPABILITY WE NOW ENJOY IS THE RESULT OF MANY OUTSTANDING SCIENTIFIC DISCOVERIES AND TREMENDOUS ENGINEERING ADVANCES, ALL CONDUCTED IN OUR FREE ENTERPRISE SYSTEM. AND I THINK THAT COMBINATION OF CONDITIONS WILL CONTINUE AND PREVAIL.

So, we are equipped with defined problems, an industry with a proven track record, and some exciting challenges. And a lot of reports - there is the Lax Report, the FCCSET (pronounced FICK-SET) Report, The Bardon Report, the Jipdee Report, and the GMD Report, to name a few, all complete with working groups, workshops, and studies. At the same time, industry has established several research entities to address our problems, such as the Microelectronics and Computer Technology Corporation, the Semiconductor Research Corporation, and ETA Systems, Inc. Also, we can point to the installation of three university supercomputer centers at Purdue, Colorado State, and Minnesota, all at no Federal expense.

How then, are we going to solve these problems of supercomputer access and waning capability? They are problems which are national in scope, so there must be interest at the Federal level. Then, how far do we carry this Federal interest? In the not-too-distant past we have tended to smother problems with money, hoping the guidance would follow.

Well, we do have some guidelines. The computer industry has responded to perceived and stated needs. Not one cent of Federal funds was expended in the research and development effort. The Federal role was to be the customer for the first of a kind.

THAT DID REQUIRE A SUBSTANTIAL AMOUNT OF GOVERNMENT ACTIVITY: SPECIFYING REQUIREMENTS, WRITING SUPPLEMENTAL SOFTWARE, AND SHARING EXPERIENCES WITH SUPPLIERS. IN TURN, INDUSTRY DEDICATED LARGE EFFORTS TO FURTHER RESEARCH AND DEVELOPMENT. UNFORTUNATELY, IT APPEARS THAT WE DID NOT CALL ON OUR COLLEGES AND UNIVERSITIES TO PARTICIPATE TO THE EXTENT OF THEIR CAPABILITIES: IN TRAINING SCIENTISTS, DESIGNING SOFTWARE, OPERATING FACILITIES, AND CONDUCTING RESEARCH. WE SHOULD KEEP ALL THESE THINGS IN MIND AS WE MOVE AHEAD.

MR. CHAIRMAN, I LOOK FORWARD TO HEARING FROM OUR WITNESSES. HOPEFULLY WE WILL OBTAIN SOME SOLUTIONS TO THESE PROBLEMS.

THANK YOU.

Mr. **Winn**. Mr. Chairman, I haven't had a chance to talk to Mr. Skeen yet. I have been handed an opening statement. I don't know if he wants to put it in the record. I will yield to the gentleman from New Mexico.

Mr. **Skeen**. I thank the gentleman, and if I could have unanimous consent just to have the statement entered into the record.

Mr. **Fuqua**. Without objection, so ordered.

[The opening statement of Mr. Skeen follows:]

HONORABLE JOE SKEEN

OPENING STATEMENT
HEARING ON SUPERCOMPUTERS
NOVEMBER 16, 1983

THANK YOU, MR. CHAIRMAN.

WITHOUT OBJECTION, I WOULD LIKE TO ENTER INTO THE HEARING RECORD A BRIEF SUMMARY OF A RECENT CONFERENCE HELD AT LOS ALAMOS NATIONAL LABORATORY ENTITLED, "FRONTIERS OF SUPERCOMPUTING."

THE FIVE-DAY CONFERENCE WAS ATTENDED BY STAFF, AND I UNDERSTAND IT WAS EXTREMELY HELPFUL IN QUANTIFYING SOME OF THE IMPORTANT ISSUES TO BE DISCUSSED THIS MORNING.

Los Alamos
Los Alamos National Laboratory
Los Alamos, New Mexico 87545

Computing Division

September 8, 1983
C-DO/BB-92-83-3369A

Dear Colleague:

For your information, enclosed is a short summary of the Frontiers of Supercomputing meeting. This item appeared in the September issue of the Los Alamos Computing Division Newsletter. A more detailed review of the Conference will appear in the next issue of Los Alamos Science, a copy of which will also be sent to each of you.

The Frontiers Conference continues to draw high praise. Thank you for your participation.

Best regards,

Bill Buzbee
Assistant Division Leader for Research
MS B260

BB:cp

Enc: a/s

Cy: CRM-4 (2), w/o enc., MS A150
Frontiers File
File

Conference on the Frontiers of Supercomputing

Motivation

The United States has been the world leader in supercomputer technology for the past thirty years, but this leadership is now being challenged by foreign competition, especially by the Japanese. Continued leadership is important to U.S. economic and national security interests, because supercomputers are a vital tool in advancing the frontiers of science and technology. To help maintain U.S. leadership, the Los Alamos National Laboratory and the National Security Agency jointly sponsored a conference entitled *Frontiers of Supercomputing*, which brought together researchers, executives, and policy makers from universities, industry, national laboratories, and the federal government.

Specific sessions at the conference covered national and industry perspectives; advanced architectures; supercomputers at Los Alamos National Laboratory; software, algorithms and applications for supercomputers; and national priorities for supercomputers. Speakers and attendees included Donald Kerr, Director of Los Alamos National Laboratory; Jeff Bingaman, U.S. Senator from New Mexico; Gen. Lincoln Faurer, Director of NSA; Kenneth Wilson, Nobel Laureate; William Norris, Chariman of Control Data Corporation; John Rollwagen, President of Cray Research; B. R. Inman, President of Microelectronics and Computer Technology Corporation; Richard DeLauer, Under Secretary of Defense for Research and Engineering; Robert S. Cooper, Director of DARPA; and many others selected for their contributions and influence in supercomputer research, development, and applications.

The Conference

Donald Kerr noted in his opening remarks that there are actually two frontiers of supercomputing — an intellectual frontier that is primarily technological and a national leadership frontier that must discover new policies, new organizations, and new collaborations.

The Leadership Frontier

Jeff Bingaman raised the question of what sort of industrial policies the country should have and the proper role of government therein. Senator Bingaman noted that Congress needs to be more aware of computing technology and of its importance.

Challenges to be Faced

B. R. Inman, in his keynote address, stated that the United States faces two challenges that must be balanced during the remainder of this century. The first challenge is military and primarily concerns the U.S.S.R. We must be prepared to deal with steadily increasing conventional forces that are highly mobile. Inman also noted an increased willingness to use this force as evidenced by recent events in Angola and Ethiopia. In subsequent remarks, DeLauer and Cooper stated that computing technology would play a key defense role "in the dangerous years that lie ahead."

The second challenge is an economic one, and centers around the third world. We must encourage political stability and make it possible for the third world to share in economic prosperity. Mr. Inman believes that the information industry will play a key role in future economic prosperity and that the U.S. should enjoy a large share of it. However, to do so, we must have properly trained people, and this has certain implications for educational systems, e.g., more math, science, and languages.

Obstacles to be Overcome

William Norris, in his banquet address, stated that the cost of American companies duplicating each other's basic research is significant and that ways are needed for American companies to share, instead of duplicate that cost. He noted, however, that such a move has both organizational and legal implications. Mr. Norris also announced the formation of a new company whose charter is to build supercomputers.

Kenneth Wilson emphasized that it is crucial for academia to participate in the utilization and development of supercomputing technology. First, we need to educate the supercomputer architects and developers of the future. Second, because so few students have been trained in and exposed to supercomputing technology, the technical and administrative staffs of industry are generally unaware of supercomputer capabilities. Yet computer simulation is often the most cost effective and rapid method for developing new products and processes.

John Rollwagen reported that Cray Research has considered ways to place supercomputing equipment in academia on favorable terms. However, unfavorable tax considerations have prevented them from doing so. Tax incentives would be promising mechanisms for nurturing the advancement of supercomputer technology.

Planning for Leadership

The Federal Coordinating Committee for Science and Engineering Technology has asked the DOE to prepare a national plan to maintain U.S. leadership in supercomputers. James Decker, DOE Headquarters, reviewed the current status of the plan. It includes provisions for accelerating government usage of supercomputers, government acquisition and usage of experimental systems, R&D tax incentives for industry and government support for long range R&D.

Very-large-scale-integration componentry (VLSI) will provide the building blocks of future supercomputers. Several speakers indicated that American manufacturers are buying an increasing amount of VLSI componentry from Japan. As noted by Jack Worlton, Los Alamos, the absence of high-speed componentry manufactured in this country that is suitable for usage in supercomputers constitutes an Achilles heel.

The Technical Frontier

The trend toward parallel computing architectures was evident in presentations by both computer manufacturers and academic researchers. By the end of the decade, commercially supplied systems with up to eight processors will be available. Effective use of parallel architecture hinges on development of algorithms and software which, when combined into a system with parallel architecture, achieves high performance. Therein lies a rich collection of research topics.

Most academic projects make extensive use of VLSI, exploiting 32-bit microprocessors and 256K memory chips. Many of them can be classified as "danceball" machines. Danceball systems have processors aligned along one side, memories on the other, and a communication geometry to bind them together. High performance on these systems necessitates algorithms that keep all of the processors "dancing" all of the time.

Three General Models

Professor James Browne of the University of Texas observed that there are three general models of massively parallel systems based on whether decisions concerning the communication, synchronization, and granularity of parallel operations are bound at design time, compile time, or execution time.

Examples of systems where decisions are made at design time are fixed-net architectures such as systolic arrays. Such systems provide fast execution of the algorithms for which they are designed, but other algorithms may have to be adapted to fit the architecture.

In bind-at-compile-time architectures, the topology of the interconnection of processors and memory can be flexibly reconfigured to suit a particular algorithm, but it will remain in a fixed configuration until explicitly reconfigured for the next problem. The University of Texas' TRAC is an example of bind-at-compile-time architecture. To allow reconfiguration, these architectures must have more complex control systems than those that bind at design time, and the reconfiguration process does take time. The trade-off is that reconfigurable architectures can be used for many different algorithms.

Bind-at-execution-time architectures do not explicitly specify topology but, through use of control structures and shared memory, allow any topology.

These systems are the most flexible since no algorithm is constrained by an inappropriate architecture. The price is the overhead required to synchronize and communicate information. Dataflow machines are an example of bind-at-execution-time systems.

Many Issues to be Resolved

Today, effective use of a multiprocessor system (and a vectorizing system) requires intimate knowledge of the hardware, the compiler's mapping of code onto the hardware, and the problem at hand. On systems with 100 or more processors, software must free the user from these details because their complexity will often defy human management. Many issues involving algorithms, models, architectures, and languages for massively parallel systems have yet to be resolved.

John Armstrong of IBM, in his discussion of VLSI technology, indicated that as the level of integration increases in componentry, there is an increased need for basic research and development in material science and packaging technology. Furthermore, he believes the current level of research in these areas is inadequate.

Several speakers noted that supercomputers systems require highspeed peripherals and that, in general, little work is being done to advance the performance of these devices. Resolution of this problem will likely require collaboration between manufacturers and government laboratories.

Traditional applications for supercomputers include nuclear weapons, aeronautical design, and fusion energy studies. Presentations from these areas showed a need for supercomputers up to 200 times faster than state-of-the-art equipment. Other presentations revealed potentially new application areas such as VLSI design, CAD/CAM, and human behavior studies. Overall, there was a high level of confidence among conference participants that supercomputer usage is expanding.

Summary

Implicit in this conference was an array of national strengths — an entrepreneurial supercomputer industry, robust and innovative academic research in parallel computation, and a growing set of supercomputer application areas. Explicit in the conference was the realization that U.S. leadership in supercomputing is important and that continuation of it necessitates a community of action between industry, academia, and government laboratories. Ultimately, the success of this conference hinges on the success of that community. In the meantime, there was wide agreement that the conference made an important contribution towards that end.

Bill Buzbee, 7-1449
Computing Division Office

Mr. **Skeen**. I thank the chairman and thank the ranking member.

Mr. **Winn**. Thank you.

Mr. **Fuqua**. Without objection permission will be granted to take photographs, recordings, or television retroactively in the hearing today.

Mr. **Fuqua**. Dr. Knapp, we are pleased to have you, and you may proceed.

STATEMENT OF DR. EDWARD KNAPP, DIRECTOR, NATIONAL SCIENCE FOUNDATION

Dr. **Knapp**. Thank you, Mr. Chairman.

I am pleased to appear before you to discuss the issue of academic access to supercomputers. With your permission, Mr. Chairman, I will submit my prepared statement for the record and proceed with a few summary remarks.

Mr. **Fuqua**. Yes; without objection we will make it a part of the record.

Dr. **Knapp**. I will begin by describing the problem of limited supercomputer access by academic scientists and the effect this has had on research and training at colleges and universities in the United States. I will close by describing the ways in which the National Science Foundation is addressing this problem.

During the 1950's and 1960's computers revolutionized research in almost every field of science. At that time the Federal Government funded state-of-the-art computers at universities and colleges. However, in the seventies Federal Government support for academic computing slackened and academic computing facilities were unable to keep up with the whirlwind of development in computing technology: minicomputers, microcomputers, computerized instrumentation and supercomputers.

In recent years NSF has supported medium-scale computers and computerized instrumentation as part of its increased support of instrumentation. This was not the case with supercomputers, however.

Academic scientists soon fell behind their industrial and government laboratory counterparts in obtaining and using the supercomputers that were needed for continued advances in such diverse areas of study as materials research, applied mathematics, physics, economics and civil engineering.

However, leading edge research in these and other fields contains a component dependent upon large-scale computation and the academic research community was effectively cut out of competition for research advances in these areas.

Some university-based scientists obtained access to large-scale computing facilities at the national laboratories or other Government installations if their fields of science overlapped with scientists in the laboratory. Even so, this definitely was the exception rather than the rule.

In addition, it was difficult to involve graduate students and postdoctoral researchers in projects that involved supercomputers. A major strength of American science has been the training of graduate students in frontline research, from both the educational value of the research done and from the value of the new ideas young graduate students bring to research environments. This major input has been lacking in fields using large-scale computing facilities and American science has suffered as a consequence.

Let me give you some examples of scientific problems that require large-scale computing needs. In materials research, scientists are solving such complicated problems as the electronic structure of complex bulk solids, polymer and biomaterials utilizing large-scale supercomputer simulations. Future applications of these studies directly affect solid state electronic technology, among other things.

In biology, scientists have set up a data base for DNA and RNA sequences. This data base, located at Los Alamos, already contains over 2 million sequences and is being used to help unravel the mystery of heredity and cell function.

In economics, scientists have used supercomputers to model macroeconomic problems. These economists look forward to expanding their dynamic nonlinear modeling capabilities as more supercomputer time becomes available, studying such things as inflation and economic growth.

Because of our close ties to the academic community, NSF recognized the problems of limited access to supercomputers and formed a panel in cooperation with other Government agencies on large-scale computing in science and engineering. This panel, headed by Dr. Peter Lax, worked to determine the scope of the problem and suggest possible actions to alleviate it. In December of 1982 NSF issued the report of this panel, which urged the Federal Government to take measures to correct this problem.

In April of 1983, under the direction of Dr. Marcel Bardon, NSF organized an internal working group on computers for research to help develop an appropriate course of action for the foundation to meet the computing needs of academic science and engineering.

Building on their observations and recommendations, an NSF task force on advanced scientific computing resources, under the direction of Dr. Edward Hayes, is developing a program aimed at making supercomputers accessible to scientific and engineering researchers. This is being done in cooperation with other concerned Federal agencies and involves short-term and long-term plans.

A key question in both short-term and long-term plans is how much supercomputer time academic researchers in this country need. Evidence provided to us suggests that considerably more time is needed than is now available.

NSF is now gathering more information from its current awardees whose research would benefit from access to supercomputers. Accurate estimates of needed supercomputer services are difficult to obtain. In many branches of science, researchers who have not had access to such facilities simply don't know the power that such facilities can bring to their research. Because of this, they can't accurately judge their supercomputer needs.

In my own specialty, particle accelerator design, we used supercomputers to make accelerator design simulations which could not be done in any other way. We developed many new approaches to that technology at very low cost. Almost all fields of science can benefit from such improved simulations that are possible only with supercomputers.

In the near term, to judge more accurately the needs, workshops and other meetings are being organized on a disciplinary basis to explore the specific supercomputer needs of the scientific and engineering community.

While these things are being determined, we are planning to take advantage of sites that can provide blocks of class VI computer services in fiscal years 1984 and 1985. University scientists and engineers are being encouraged to submit proposals for class VI computer services and appropriate ancillary costs such as travel, communications, and local equipment.

All of this, of course, will take time. Before there can be a major expansion of research activities with supercomputers, the groundwork must be laid. This means providing money and expertise for software conversion and code development.

We have issued a solicitation for large blocks of class VI computer services to meet the near-term demands of our grantees. Preliminary returns indicate that we will get good response from both universities and private companies wishing to offer large blocks of computer services.

For the long term, the Foundation is exploring more stable arrangements for providing access to supercomputers, such as university-based supercomputing centers. The key here is long-term. Although some researchers will be able to make immediate use of supercomputers in their research, many others will need to dedicate 2 or more years of their time and energy to develop algorithms and programs before they can even begin to make effective use of large-scale computers. To make such effort worthwhile, they must have the security of long-term arrangements for access to supercomputers.

Our long-term objectives include not only providing large-scale computer resources for research and training of a new generation of computational scientists and engineers, but improving remote access to advanced computing facilities for the scientific and engineering community. We will do this by extending and developing appropriate communication systems and networks between supercomputer centers and users.

We also plan to develop advanced computational resources locally at universities. In fact, in January we are planning to establish an NSF program advisory committee to advise the Foundation on policies and issues related to supercomputer systems at universities. Prominent among the issues to be discussed will be the training of the people necessary for such systems to work effectively.

In a nutshell, we will build on existing computer centers, especially those at universities. At the same time, new ones will be considered and established when needed by this country's scientific and engineering community.

An advisory committee will be established to consider such things as research needs and ways to support computer resources, the merit of research projects, allocation of computer services, and communication, and networking needs.

I would like to emphasize that our work has been and will continue to be coordinated with the many active and potential participants in the use of supercomputers and in computer-related research. These participants include Federal agencies, existing university supercomputer centers, computer manufacturers, and other private companies, and nonprofit corporations.

In addition to our activities aimed at improving access to supercomputers for research, I would like to note that we will continue to support research in the theoretical and experimental design of computer systems and in computational mathematics, software, and algorithms.

This concludes my statement. I will be pleased to answer questions that you may have.

Mr. FUQUA. Thank you, Dr. Knapp. We will hear from Dr. Martin and then we will have some questions for you.

[The biographical sketch and prepared statement of Dr. Knapp follow:]

NATIONAL SCIENCE FOUNDATION
WASHINGTON D C 20550

DR. EDWARD A. KNAPP
Director

Dr. Edward A. Knapp became Director of the National Science Foundation in a recess appointment announced on November 2, 1982, by President Reagan. He had been Assistant Director for Mathematical and Physical Sciences of the NSF since September 16, 1982. Prior to coming to the NSF, he was head of Los Alamos National Laboratory's Accelerator Technology Division in Los Alamos, New Mexico.

Dr. Knapp was born in Salem, Oregon, on March 7, 1932. He earned his A.B. degree in physics in 1954 from Pomona College, and his Ph.D. degree in high energy physics from the University of California (Berkeley) in 1958. He joined Los Alamos National Laboratory in 1958, and was involved with the pioneers of the Manhattan Project in the early work on controlled thermonuclear research. His group made significant contributions to the understanding of energy loss mechanisms in plasma discharges.

Dr. Knapp participated in the initial planning for the Laboratory's half-mile-long linear proton accelerator, the Los Alamos Meson Physics Facility (LAMPF). While LAMPF was being constructed, Dr. Knapp specialized in accelerator physics and headed the research group that designed the accelerator structures for the facility. These systems were also utilized in industry, and ultimately led to a compact electron accelerator that is now used by more than 500 institutions for radiation therapy of cancer.

In 1972, Dr. Knapp spent a sabbatical at the European Organization for Nuclear Research (CERN) in Switzerland, where he worked on the design of a new proton accelerator for that facility and the development of a collaborative high-energy physics program involving Los Alamos and CERN. On returning to Los Alamos in 1973, he assumed responsibility for the applied physics program and for the medical pion beam at LAMPF that was used successfully for several years for the treatment of several types of cancer.

Dr. Knapp is a Fellow of the American Physical Society, a member of many scientific and technical organizations, and is the author or co-author of almost 100 scientific papers. He has served on many panels and committees in his fields of interest.

November 1982

STATEMENT OF
DR. EDWARD A. KNAPP
DIRECTOR
NATIONAL SCIENCE FOUNDATION

BEFORE THE

COMMITTEE ON SCIENCE AND TECHNOLOGY
HOUSE OF REPRESENTATIVES

NOVEMBER 16, 1983

Mr. Chairman and Members of the Committee:

I am pleased to have this opportunity to meet with the Committee on Science and Technology during its hearings on supercomputer design, development and use. I wish to emphasize at the outset that needs for supercomputers arise in many scientific and engineering disciplines. I am here today to speak to the issue of the needs for supercomputers in all scientific and engineering disciplines supported by the National Science Foundation.

Computing facilities can have an important effect on research which is conducted by academic scientists and engineers. During the 1950's and 1960's, the Government encouraged the growth of computing in academic research; research methods were transformed in many disciplines; and the U.S. enjoyed a large, lead in quantitative research and modeling complex phenomena.

During the period of the 1970's, although minicomputers, microcomputers, automated measuring instruments and a variety of other computerdependent tools were introduced into academic laboratories, academic researchers appeared to lag behind their counterparts in industry and at National Laboratories. However, in two disciplinary areas -- atmospheric and ocean sciences, and magnetic fusion energy sciences -- national research computing facilities, were established.

I am personally familiar with the National Magnetic Fusion Energy Computing Center supported by the Department of Energy (DOE) and with other DOE computing centers; however, others will describe these centers to you. I would like to tell you briefly about the National Center for Atmospheric Research (NCAR) at Boulder, Colorado. NCAR has been providing computational services to the atmospheric science community for twenty years. The acquisition of successive generations of computers has made it possible for scientists to address problems in the theory of weather prediction, climate studies, severe storms research, ocean circulations, and the influences of the sun, land, and oceans on the atmosphere. Today, NCAR provides computational resources to more than 900 scientists at 75 institutions. In addition to NCAR's important role in serving the computational needs of the atmospheric science community, NCAR has become a leading computational resource for the ocean sciences.

Within the overall Federal role for research in science and engineering, the National Science Foundation has the lead responsibility for basic research and science education conducted in academic institutions. By the National Science Foundation Act of 1950, as amended, the Foundation is authorized and directed to

> "...initiate and support basic scientific research and programs to strengthen scientific research potential and scientific education programs...in the mathematical, physical, medical, biological, engineering, social, and other sciences...
>
> and
>
> "...foster and support the development and use of computer and other scientific methods and technologies, primarily for research and education in the sciences..."

Thus, from both a general and a specific perspective, the Foundation has a mandate to assume a major responsibility with respect to computing for research and education in the sciences and engineering.

Historically, the Foundation has interpreted this responsibility to include programs designed to promote access to and improve the capacity and capability of whatever computing facilities may be necessary for the solution of fundamental problems in basic research as well as education in the use of such facilities.

In April 1983, I established a Working Group on Computers for Research under the direction of Dr. Marcel Bardon to assess needs of the science and engineering community for access to advanced large-scale computers. The Working Group issued a report entitled A National Computing Environment for Academic Research. It is appended for the Record.

At the outset it is important to emphasize two aspects of NSF's activities regarding supercomputers: first, the coordination of efforts among the various government agencies, private concerns, universities and other organizations is essential to ensure effective utilization of human and other resources in support of scientific and engineering research and education. Second, for the Foundation, a key element will be the use of advisory committees and peer review. Past experience indicates that the advice of discipline advisory committees is very important in maintaining the quality of scientific and engineering research and education.

In this statement, I will address three vital considerations about supercomputers in science and engineering:

- Scientific and engineering research needs for supercomputers
- Research on supercomputers
- What NSF is doing

1. SCIENTIFIC AND ENGINEERING RESEARCH NEEDS FOR SUPERCOMPUTERS

The scientific need for greater access to supercomputers has been addressed in several studies [1, 2, 3,] and the progress in research that can be made using computers has been the subject of many publications. In May 1983, the American Institute of Physics devoted an entire issue of its publication entitled Physics Today to the subject of "Doing Physics with Computers" [4]. In this issue, applications of large-scale computers to surface physics, high energy physics, condensed matter physics, astrophysics, plasma physics, atmospheric physics, and other fields are described.

As part of the NSF staff study [1] several types of problems were identified that currently or may in the future require large-scale computation for solution. Some of the problem areas are:

Astronomy - Understanding convective stellar interiors and relativistic interstellar jets involves hydrodynamics, magnetohydrodynamics, radiative transport, atomic physics, and an understanding of the transition from order to chaotic behavior. These are highly nonlinear systems which cannot be solved analytically but can be studied by numerical simulation.

Chemistry - Simulations are valuable in the exploration of many-body effects that occur in dense gases. For example, simulations of three-body recombination helps develop understanding of chemical reactions such as combustion.

Materials Research - Problems requiring large-scale computing include the electronic structure of complex bulk solids, polymers and biomaterials, impurities and vacancy defects, disordered solids, and amorphous materials such as modulated structures and sandwiches. Several of these problems focus on the unexplored regime which lies between individual molecules and solid aggregates which is important for the future development of electronics.

Physics - In continuum quantum field theory, there is no known method for numerical approximation of the non-perturbative regime. However, the formulation of lattice versions of currently popular gauge theories has led to substantial new insights into the theory of quark confinement, and has also opened up the application of statistical mechanical high temperature expansion techniques to study the behavior of gauge theories for large values of the gauge coupling constant.

Civil Engineering - Advanced large-scale computing facilities are required for studies of flows and the effect of pollutants in underground aquifers, evaluation of coastal erosion, behavior of large complex structural systems, simulation of boundary layer wind flows and their interaction with civil engineering structures, and the numerical analysis of large complex problems using finite element techniques such as the analysis of material behavior.

Chemical and Process Engineering - Mathematical models of chemical processes and the prediction of chemical and physical properties of interest to chemical engineers increasingly must involve nonlinear partial differential and integral equations in several space variables and in complicated geometries. The solution of these complex models requires the techniques of modern computational mathematics and the computing power of current and future large-scale scientific computers.

Electrical Engineering - The design of Very Large Scale Integrated (VLSI) circuits has become so complex and the fabrication of prototype circuits so expensive that simulation of the circuit behavior using large-scale computers is necessary to assure internal consistency of the design before building and testing it.

Applied Mathematics - Modeling of the fluid dynamics of blood flow through the valve system of the heart has recently led to a much improved understanding of how the valves work. This is a highly non-linear system which can only be studied with numerical simulation techniques.

Biology - A data base for DNA and RNA sequences has been set up at Los Alamos which contains about 2,000,000 sequences. These sequences are of variable length, up to 1000 residues long, and to search this data base for sequences "close" to a given sequence requires about 10 minutes of Class VI computer time.

Economics - The rigor and size of dynamic nonlinear economic modeling problems require, in many cases, Class VI computer capability.

The underlying mathematical and computational problems that arise in several of these examples are very similar in that the solution of nonlinear partial differential equations by numerical simulation is required. Combinatorial problems with many degrees of freedom arise in other areas. Since the algorithms that must be developed and programming problems that must be solved by scientists in these diverse disciplines have much in common, there is substantial benefit to be gained from increasing the communication among scientists and engineers in different disciplines. An important aspect of NSF's planning is the support of computer networks for academic research not only because they make access to supercomputer facilities easier, but also because they promote and facilitate communication among scientists with common computational interests.

Many of the problems that require supercomputers for their solution involve simulations of the behavior of complex physical, biological, or social systems. This use of large-scale computing in research highlights one of the unique contributions of computational methods. In many cases it is not possible to isolate the components of complex systems in experiments to understand the influence of each on the phenomenon being studied. One must substitute, instead, a theoretical exploration of the phenomenon itself based upon knowledge of the underlying scientific principles. This leads to complicated, nonlinear problems in which researchers conduct numerical experiments instead of laboratory experiments. In fact, in some cases the laboratory experiments cannot be performed because they would be prohibitive in cost, dangerous or not sufficiently well controlled to yield useful data.

An example of the role of large-scale computing in science and engineering is the work of astronomer, Professor Larry Smarr, University of Illinois. Professor Smarr is interested in explaining the origin and behavior of gas jets, lobes and hot-spots that are being observed by radio astronomers using the high resolution of the Very Large Array (VLA) radio telescope. Controlled experiments are impossible and laboratory experiments offer limited promise because the cost of studying relativistic jets in the laboratory would be prohibitive even if the technical problems involved were understood. As a result, Smarr and his colleagues have undertaken a program to study basic jet physics numerically. This is a study that requires hundreds of hours per year on a supercomputer over several years time.

Professor Smarr is collaborating with scientists at Max Planck Institute in Garching, West Germany, where he has access to the supercomputer time and computer graphics facilities. Other scientists in this and other fields are also being offered access to supercomputers outside the United States.

Three universities currently have supercomputers. User communities and the underlying financial bases are being developed incuding much emphasis on expanding beyond the local campus. It is unlikely, however, that such systems will be commonplace on many university campuses in view of the capital and operating costs required relative to the need that would be generated within a single university. Supercomputers should, therefore, be treated as scarce resources. They will ordinarily have to be shared by researchers in diverse geographical locations and possibly a wide range of research disciplines. This situation only underscores the need for more effective communication among scientists and engineers and between these researchers and the computing resources utilized in their research.

Researchers also stress the need for access to computers that are vastly more powerful than the supercomputers available today or in the near future. Factors of 1000 have been mentioned as a future goal. While factors of 10 are feasible with technological improvements, it is not known how to build a machine that is 1000 times faster than those currently available nor how to develop problem strategies or algorithms for them. The power or capability of a computer is important in this regard because the elapsed time, i.e., researchers waiting time, for computations to be performed by the computer must not be so long that the research becomes infeasible to do. Problems which need much more powerful computers for solution include:

- Biology - The simulation of the structure of a bi-layer membrane in contact with ions in water would permit tracing the possible paths determined by polypeptides for ion transport throught the membrane. It is estimated that much greater computational power may be needed for this fundamental problem of biological significance.

- Chemistry - Chemical reactions can be modeled using the basic equations of quantum mechanics. To model reactions involving larger, more interesting, molecules substantially more computational power is needed.

- Materials Research and Electrical Engineering - As dimensions get small in microelectronics, classical physics fails. Quantum mechanics is needed to model small-dimension physics for semiconductor devices.

- Atmospheric and Ocean Sciences - Calculations for atmospheric modeling require finer resolution of grid points and longer time scales; modeling ocean basins would require about 100 times the current capability.

- Earth Sciences - Modeling dynamic magnetic systems, for example, to determine why the Earth's magnetic field reverses--requires capability beyond that currently available.

- Astronomy and Physics - Astrophysical problems such as modeling the processes by which stars are born require much greater capability than current computers offer.

- Psychology and Information Sciences - Current computers are not sufficiently powerful to model a human's complete knowledge, even of one limited subject area.

- Economics - Inflation, unemployment, economic growth, international trade balances and other important economic aggregates result from decision-making by millions of individuals. To completely model this dynamic process is beyond the capability of existing computers.

2. RESEARCH ON SUPERCOMPUTERS

Research underlying radical new computer architectures required for the scientific problems envisaged in the future as well as research on computational mathematics, new problem-solving strategies, new software systems and experimental testing and evaluation of design concepts for numerical computing is conducted primarily in universities. It is also important that a new generation of scientists and engineers be trained in computational mathematics and computer techniques to respond to the future research needs.

It is my understanding that the long range goal of the DARPA Program in strategic computing is to develop information systems (hardware and software) which will support a range of military applications well beyond the state-of-the-art. These applications include autonomous systems, real-time battlefield assessment, knowledge-based military planning and advising, and the simulation of very large, complex systems. In addition to the military applications, the program will develop generic applications such as speech, vision and natural language. These objectives are distinct from those of the Foundation. While the DARPA effort may contribute to new computer architectures, the effort does not address the needs for advances in numerical computations characterizing scientific and engineering research. Neither does it specifically address the research in computational mathematics and applications software important to the effective and efficient use of such computers in research and particularly the training of graduate students in these areas. Finally the DARPA effort is not directed at research on networking linking researchers to the supercomputers of the future.

Several years ago NSF recognized the existence of a problem with the infrastructure of computer research. Academic research in computer science and computer engineering was heavily oriented to theoretical areas in some cases because universities lacked the equipment and support staff to conduct experimental research. In response to this problem the Foundation started its Coordinated Experimental Research (CER) Program in FY 1980. This is a program of awards to institutions which have strong research capability but lack some ingredients for experimental research. Each award is designed to provide research equipment, support staff and other costs necessary to give these institutions the capability to conduct experimental computer research. To date the Foundation has made fourteen such awards.

This program has been successful in strengthening the quality of academic computer research. Universities that have received these awards have been better able to recruit and retain high quality faculty and students, have increased the quality and quantity of their research, and have become much better able to develop support from other sources, both private and governmental.

The University of Washington, which was the recipient of the first of these grants, provides a good example. They are now in the fourth year of a CER grant. The number and quality of research activities has increased substantially during the course of the grant. The number of faculty has increased approximately twenty percent, the number of Ph.D. graduate students by a similar amount. The new research activities of the university have attracted the attention of several companies; Intel, Digital Equipment Corp., Xerox, and Boeing Aircraft Corp have all formed close relationships with University of Washington research groups. Also, other federal agencies have been attracted by the new capabilities of the university researchers and have increased their support. As a result, the department has more than matched the support provided by NSF with other support from private and other federal sources. One of the new activities of great promise which has started since the grant was awarded, is work on new, highly parallel computer architectures and Very Large Scale Integrated circuit design, work which is important for supercomputers of the future.

3. WHAT THE NATIONAL SCIENCE FOUNDATION IS DOING

The recommendations of the NSF Working Group on Computers for Research [1] are important elements of the Foundation's planning to provide high-quality access to advanced computational facilities to the academic research community. The planning is Foundation-wide in support of all computationally-intensive disciplines.

In response to these recommendations, I established in September of this year, an NSF Task Force on Advanced Scientific Computing Resources under the direction of Dr. Edward F. Hayes to develop Foundation plans for providing access to advanced large-scale computing resources for the scientific and engineering community. The Task Force is developing near-term plans for providing access to supercomputers (Class VI computers) for researchers during Fiscal Year 1984, as well as long-term plans.

The key question, of course, is how much supercomputer time is needed for academic research in the U.S. The evidence provided by some scientists and engineers appears to suggest that more is needed than is now available. The Foundation is querying its current awardees to elicit information from those whose research would benefit from access to supercomputers. In addition, workshops and other meetings are being organized on a disciplinary basis to bring together researchers to explore in detail supercomputer needs of the scientific and engineering community.

The near-term plan is to take advantage of those sites that can provide blocks of Class VI computer services in Fiscal Years 1984 and 1985 which will then be allocated to the most meritorous projects where researchers have requested such computer services. A Project Solicitation has been issued to request proposals from institutions and organizations that can provide Class VI computer services to researchers. Concurrently, NSF-supported scientists and engineers and other researchers are being encouraged to submit proposals for Class VI computer services and appropriate ancillary costs, e.g., local equipment, travel, communication, where their research will benefit from such services. In some instances, it is expected that these latter proposals will contain two phases, a preliminary phase for software conversion and code development for use of a Class VI computer followed by a longer-term research phase. The Task Force will coordinate the review of proposals received.

At the same time, the Foundation is seeking to improve access to large-scale computing for academic research on a long-term basis and is preparing plans for future years to achieve this goal. A longterm commitment is essential to the success of this program because researchers, in many cases, would need to dedicate two or more years of their time and energy developing algorithms and programs before they could make effective use of large-scale computer resources.

Objectives in the long-term include:

- providing the large-scale computer resources to meet the needs of U.S. scientific and engineering research including development of the new generation of computational scientists and engineers;
- improving remote access to advanced computing facilities for the scientific and engineering community through extending and developing appropriate communication systems and networks;
- developing advanced computational resources locally at universities;
- supporting academic research in the areas of computer architecture, computational mathematics, software and algorithms applicable to supercomputers.

Specifically, supercomputer centers, fully configured and appropriately staffed to provide computing services to local and remote users, will be provided by building on existing centers and establishing new ones as appropriate to advance U.S. scientific and engineering research. The allocation of supercomputer resources should be based on scientific and engineering merit of research projects, the need for the speed and memory capacity of supercomputers and the efficiency of computer programs to be used. To advise the Foundation on the issues and considerations involved, e.g., research needs and ways to support computer resources, merit of research projects, allocation of computer services, communications and networks, etc., NSF intends to establish an advisory committee on advanced scientific computing.

Coordination is essential because there are many active and potential participants in the use of supercomputers and in computer-related research including federal agencies, existing university supercomputer centers, computer manufacturers and other private companies, and nonprofit corporations.

Adequate local facilities are essential to enhance scientific and engineering research and education. Almost all scientists and engineers need a wide range of appropriate computer services as well as access to supercomputers. Researchers need to develop programs, manipulate data, perform mathematical and statistical computations, communicate with collaborators and carry out many other essential professional tasks at or within close proximity of their everyday working environment. Academic scientists and engineers do not, in general, now have access to well-developed interactive local computer facilities that they could find particularly useful.

Computer networks are an important element in any long-term plan. While occasional trips to a facility may be necessary and desirable, in most cases travel to a distant site is an inefficient use of time as well as a barrier particularly to graduate student involvement in research. Remote access to supercomputers is a realistic option even with existing telecommunications technology and efforts must be supported to link university and laboratory researchers with each other and with supercomputer centers.

In parallel with these efforts to provide researchers with access to advanced computing facilities, NSF currently supports and will continue to support theoretical and experimental research in computer systems design and in computational mathematics, software and algorithms.

This concludes my prepared statement. I will be pleased to answer any questions you may have.

REFERENCES

---

[1] "A National Computing Environment for Academic Research," NSF staff study prepared under the direction of Marcel Bardon, Kent K. Curtis, Chairman, July 1983 (See Appendix)

[2] "Report of the Panel on Large Scale Computing in Science and Engineering," sponsored by DOD and NSF in cooperation with DOE and NASA, Peter D. Lax, Chairman, December 26, 1982

[3] "Prospectus for Computational Physics", sponsored by NSF, William H. Press, Chairman, March 15, 1981

[4] "Special Issue: Doing Physics with Computers", Physics Today, 36, 5, May 1983

Mr. FUQUA. Dr. Martin, we are very pleased to have you here this morning and glad to listen to—to hear your testimony.

STATEMENT OF DR. EDITH MARTIN DEPUTY, UNDERSECRETARY, RESEARCH AND ADVANCED TECHNOLOGY, DEPARTMENT OF DEFENSE

Dr. MARTIN. Thank you, Mr. Chairman.

I am very pleased to address this committee and describe the needs and efforts in large-scale computing under my purview. I would like to submit my prepared statement for the record and give some condensed remarks.

I will also discuss our mission critical computer hardware and software technology problems and the problems we have and the programs that we have initiated to overcome some of these problems. Their solution I feel can be expected to have a very positive impact on the supercomputer field.

It is well known that DOD has sponsored the development of the first large-scale computer in order to perform ballistic calculations back in the 1940's. Our involvement in computer technology spans that early work on to our initiation of the fabrication of integrated circuitry, up to our current hardware efforts in VHSIC and our software initiatives in Ada and STARS.

DOD science and technology priorities are driven by specific military system requirements. Today I would like to tell you where we are going in the computer area, including our use and need of supercomputers and our requirements and initiatives for high-speed mission critical computer systems. These programs are coordinated closely with DARPA's strategic computing program, which has already been presented to this committee.

DOD currently has a relatively small number of large-scale computing machines, such as the CRAY-1 and the CYBER-205 class machines. I have provided in my prepared statement a list of these machines, including the organizations which have them and their major applications. I will pick out a few here just for bringing you up to date.

Certainly we have some that address nuclear weapons effects analysis, and that's at the Air Force Weapons Laboratory; wind tunnel data reduction and aerodynamics analysis at the Air Force Arnold Engineering Development Center; atmospheric research and weather prediction at the Naval Fleet Numeric Oceanography Center——

Mr. FUQUA. Is that in Monterey?

Dr. MARTIN. Yes. Communications security and foreign intelligence at NSA, and aerospace calculations and orbital mechanics type problems. A variety of signal processing problems are also addressed by these machines.

Now, I have also described in the formal statement plans for future procurements of machines of this class by other laboratories and organizations. DOD large-scale computers for the most part are engaged in classified activities but these machines, I can assure you, are very heavily utilized. Access to these computers has essentially been limited to those researchers directly engaged in work of the organizations that are in possession of the machines.

Now, I have asked my director of research and laboratory management to determine what steps we might take to make supercomputers more accessible to other researchers in the future, be they our own laboratories or universities.

Our use of large-scale computers has primarily, but not exclusively, been in scientific calculation and supportive research. Now, however, new types of applications are beginning to emerge, and these fall into such areas as artificial intelligence, graphics, complex image processing, command and control, and space-based defense.

As one might expect the President's recently announced initiative in space-based defense will perhaps present the most severe of all future military computational problems in that the real time requirements for fail-safe detection and warning, target tracking, discrimination of thousands of space objects, including decoys, locking in or detecting warheads and guiding defending weapons is potentially one of tomorrow's most formidable tasks.

The raw computational power required for such a system is far beyond that of any other system contemplated today, but let us not forget the software challenge that it is going to bring. The software for all previous systems, including the space shuttle, are going to be dwarfed by comparison.

Another sometimes overlooked large-scale computing requirement is that of simulators to drive the operational system software. They must generate artificial threats and environments and they too must run in real time.

A space-based defense system will have some characteristics typical of our high-end signal processing systems. The extremely high data rate of incoming raw signals by our modern radar and optical sensors was the very situation that motivated the establishment of the VHSIC program that we have today and was initiated in 1980.

Now let me turn to advancements in supercomputers. In the past two decades we have seen some tremendous advances which have come primarily from increased speed in the basic componentry or circuits. Such circuit advancements are becoming more difficult to achieve as we approach some of the technological limitations of the implementing technologies.

Highly parallel computer architectures are the basis for significantly greater improvements in the future. Developments of such architectures in the form of supercomputers is an extremely difficult task unto itself, but the use of a system of very large numbers of parallel processors presents an extremely demanding challenge to us. Much work on algorithms, languages, and support software is required. We do not yet know how to match these architectures with our conventional computing problems.

DOD has made some advances in distributed architectures. Now, this is a form of parallelism which is extremely important for the survivability and failsoft requirements, but efficient use of parallelism and distributed systems is also very difficult to achieve.

The department has a joint effort with the Department of Energy on a highly parallel computer system known as the S-1. It will have 16 processors, each approximating a CRAY-1's performance. These are configured in what we call a dancehall arrangement with memory down one side and the processors down the other. It is a little bit like having Democrats on one side and Republicans on the other at budget time, and we hope they can handshake and do everything that is necessary to give us the performance that is necessary. But again, exploiting this kind of parallelism is proving to be difficult and a challenge.

The department is also supporting research systolic array processing. Now these are fixed networks of relatively simple arithmetic units arrayed in configurations to match specific problems. There are no data buffers or registers between the arithmetic units. The computations are performed in a lock-step fashion and the results are immediately shifted to the next processing element.

The clocked or rhythmic nature of this processing gives rise to the term "systolic." Now while such systems can be extremely fast, I should mention emphatically that they will also be highly inflexible.

In the past, except for NSA's requirements, DOD's supercomputer needs have been somewhat less stressing than DOE's and NASA's, and DOD's needs have generally been met by the industry. Of course, we have communicated with industry in stating our requirements, so it may be a little difficult to say where the—which was the chicken and which was the egg.

Because of the highly classified nature of NSA's work, it's going to be difficult to be very specific about their requirements. They have a number of prime users. They are among the prime users of supercomputers today, and members of the agency are currently working with industry to enhance our supercomputer technology for the needs they foresee in the not-too-distant future.

Now, in all of our applications of complex computer architectures, either large or small, we face a very difficult issue, and that is matching the problems with the machines or configuring our machines to match the problems. Both of these options are open to us with new technology.

There is also the important related question, how can we create software tools to aid our efforts in managing that match? Our large-scale computer systems in the past have been delivered with little or no software tools. That I think we should come back to as one of the reasons why accessibility on the part of the academic community has been very limited.

A major portion of early—of the early program to develop new architectures must in the future also address application and software issues. If we are going to exploit our hardware advances, then they must be accompanied hand in glove with comparable advances in software capability. Software cannot be a stepchild.

Now I would like to turn for a moment to our mission critical computer requirements. They arise from more conventional tactical systems, and they derive from our policy to meet superior—to meet numerical superiority on the part of our adversary with technological superior weapons systems.

Computers represent in that vein a keystone technology for maintaining a military technological edge. Our modern technical—tactical systems demand extremely small, extremely fast and ultrareliable computer elements suitable for operation in some very severe military environments.

Our VHSIC program, which I mentioned earlier, is developing the technology to pack the equivalent of about 150,000 transistors on a single silicon chip of about one-quarter inch by a quarter inch. Phase 2 of this program will extend that technology to more than double the density and likewise provide some proportional improvements in speed.

Now, while initially aimed at signal processing requirements for tactical systems, VHSIC will benefit all forms of computing and certainly as a component will be a contribution to supercomputers of the future.

Exploitation of the fundamental benefits of digital computers and software has been accompanied by a new set of problems. One of these is the rapidly increasing demand for new, larger and more complex software application systems.

Technology to make the development and evolution of this software less labor-intensive have not kept pace with a very rapidly growing user base. Since the growth in software investment in defense is often misunderstood as increasing costs for the same capability, I would like to spend a moment to set the record straight by explaining in part how this growth is caused.

There are at least four distinct factors driving our software growth. One is rapid migration to software of functions previously performed by hard-wired devices; for example, our communications switching.

Entirely new functions impossible before the emergence of stored programs in digital devices are now being developed, and examples of that are digital map displays.

Functions previously performed by separated subsystems are being performed by a single set of time-shared components. The Air Force's integrated communications, navigation, identification, and avionics, or ICNIA, system is an excellent example of that.

The growing use of separate, digitally implemented functions operating in parallel and communicating via multiplex buses imposes new control requirements which must be implemented in software.

So, there are a lot of additional requirements being placed in software that give rise and give considerable notice to the dollars being spent in the software area in general. I think if we have criticism for the software area, it must be that the methods we use for developing and evolving that software is far too labor-intensive and that change in that dimension must be made.

We have taken some aggressive action to deal with those problems, and the first of which was our Ada program, which was initiated in 1975, and that was to achieve a common, high order language for mission critical computer systems.

Development of Ada was motivated by excessive use of assembly languages and a proliferation of poorly supported high order languages resulting in difficulties in supporting the software through the system life cycle and lack of portability of that software from one type of architecture to another type of architecture.

About 1½ years ago, after much study of what we felt were our software problems and our software opportunities, we concluded that a broader attack was needed to address the software explosion.

We called this attack the software initiative originally and later added the name STARS, or software technology for adaptable, reliable systems.

It has been suggested in the private sector that the STARS program will do for DOD and the U.S. software industry what VHSIC has and will mean to DOD and the U.S. semiconductor industry.

Now, the STARS program focuses on the creation of an integrated, highly automated development and support environment for software. This is termed a software engineering environment and it is analogous to the set of tools and capabilities known as CAD/CAM, or computer-aided design/computer-aided manufacturing for VLSI chip development and production.

The STARS program builds on the efforts initiated in the Ada program, and I've described the relationship between the two of those in my prepared statement. Now, while STARS was motivated by our needs in mission critical computer systems, it is equally applicable to large scale or supercomputing.

Our existing software tools for mission critical systems has proven to be inadequate for our growing military software demands. We are even worse off in the supercomputer area since we have had virtually no software supplied by the manufacturers. This has severely limited the application and thus the market for large-scale machines since their users essentially had to develop their own support software before achieving any product benefit.

Now, what this says is that unless one absolutely had to have the compute power of a supercomputer then one sought other solutions and other alternatives for their computer requirements.

In summary, DOD has had a long and a very deep involvement in nearly all aspects of computers and the development of hardware and the development of software algorithms. For the most part industry has been meeting the needs of conventional large-scale computing.

Large-scale numerical computing machines have been very important to DOD and will continue to be so, but they have not been the principal focus of our developments in computing. Our needs are changing and the technology certainly is advancing very rapidly. We intend to be at the forefront and be the most voracious users of the latest and greatest in computing technology, as we have been in the past.

In the not too distant future, if they can meet the environmental requirements and demonstrate the robustness needed for survivability, then we might very well see machines we now classify as supercomputers as part of mission critical computer systems.

This is a migration that has taken place quite commonplace in the past, and as we take what is known as supercomputers today and put them in our tactical systems, we will find that technology is pushing forward and there are new architectures and new frontiers being pursued by those who are in the supercomputer arena, and we intend to use the capabilities they provide as well.

And finally I would like to say that it is critical that the United States maintain its lead in supercomputer technology. DOD does benefit and will continue to benefit from progress in this area. While our programs have not been directly focused on supercomputers, we believe that they have been supportive to growth in supercomputers and I think that our initiatives demonstrate that.

Thank you.

[The biographical sketch and prepared statement of Dr. Martin follow:]

DR. EDITH W. MARTIN

Dr. Edith W. Martin is Deputy Under Secretary of Defense for Research and Advanced Technology in the Department of Defense. She was formerly a Corporate Executive with Control Data Corporation, and Director of Government Systems, Atlanta Operations. Prior to that position, Dr. Martin was Director of the Computer Science and Technology Laboratory at Georgia Tech's Engineering Experiment Station. In both of these positions she was involved in Defense-related research and development activities. Dr. Martin currently serves as the top U.S./defense official providing policy, guidance, and oversight for U.S. defense science and technology. She is the principle U.S. representative to the NATO Defense Research Group and is NAMRAD principle to The Technical Cooperative Program with our allies.

Dr. Martin has been recognized for her leadership and contributions to the defense/industrial community via numerous awards, including several from the Electronics Industries Association, the American Defense Preparedness Association, Armed Forces Communications and Electronics Association, and the Institute for Electrical and Electronics Engineers. She has also served as a member of the Defense Science Board Special Task Force on Embedded Computer Resources, on the Editorial Review Board (Computer Science Expert) of Military Electronics Countermeasures, on the Editorial Board of IEEE Software Magazine, and as chairman of a large number of technical review committees. Dr. Martin's current initiatives in technologies essential to military critical computer systems; export control and technology transfer; and science and engineering education have received strong support throughout the Department of Defense, the commercial sector and university communities. Her publications in computer science and almost the full spectrum of defense research are widely read and quoted.

Dr. Martin received her Ph.D and M.S. degrees in information and Computer Science from the Georgia Institute of Technology. She received a B.A. degree in Experimental Psychology from Lake Forest College, Illinois.

THE DEPARTMENT OF DEFENSE
STATEMENT ON
LARGE-SCALE AND MISSION-CRITICAL COMPUTING

BY

DR. EDITH W. MARTIN

DEPUTY UNDER SECRETARY OF DEFENSE

FOR

RESEARCH AND ADVANCED TECHNOLOGY

BEFORE THE

COMMITTEE ON SCIENCE AND TECHNOLOGY

UNITED STATES HOUSE OF REPRESENTATIVES

NOVEMBER 16, 1983

INTRODUCTION AND BACKGROUND

MR. CHAIRMAN, I AM VERY PLEASED TO ADDRESS THIS COMMITTEE AND DESCRIBE THE NEEDS AND EFFORTS IN LARGE SCALE COMPUTING UNDER MY PURVIEW. I WILL ALSO DISCUSS OUR MISSION-CRITICAL COMPUTER HARDWARE AND SOFTWARE TECHNOLOGY PROBLEMS AND THE PROGRAMS WE HAVE INITIATED TO OVERCOME THEM SINCE THEIR SOLUTION CAN BE EXPECTED TO HAVE A VERY POSITIVE IMPACT ON THE SUPERCOMPUTER FIELD.

IT IS WELL KNOWN THAT DOD SPONSORED THE DEVELOPMENT OF THE FIRST LARGE-SCALE COMPUTER IN ORDER TO PERFORM BALLISTIC CALCULATIONS. SINCE THAT TIME WE HAVE BEEN AND CONTINUE TO BE A MAJOR CONTRIBUTOR TO COMPUTING. OUR INVOLVEMENT SPANS EARLY WORK IN INTEGRATED CIRCUITRY UP TO OUR CURRENT HARDWARE EFFORTS IN VERY HIGH SPEED INTEGRATED CIRCUITS OR "VHSIC", AND OUR SOFTWARE INITIATIVES, ADA® AND STARS.

DOD's SCIENCE AND TECHNOLOGY PRIORITIES ARE DRIVEN BY OUR SPECIFIC MILITARY SYSTEMS REQUIREMENTS. THEY ALSO TAKE INTO ACCOUNT THE CAPABILITIES BEING MADE AVAILABLE BY THE PRIVATE SECTOR. OUR NATION'S COMPETITIVE DRIVE, MOTIVATED BY COMMERCIAL OPPORTUNITY, FOSTERS MANY IMPORTANT ADVANCEMENTS IN TECHNOLOGY. IT IS THE VITALITY OF THIS TECHNOLOGY BASE THAT PROVIDES THE U.S. WITH THE ABILITY TO STAY TECHNOLOGICALLY AHEAD OF OUR ADVERSARIES.

TODAY, I WOULD LIKE TO TELL YOU WHERE WE ARE GOING IN THE COMPUTER AREA INCLUDING OUR USE OF AND NEED FOR SUPERCOMPUTERS, OUR

REQUIREMENTS AND INITIATIVES FOR HIGH SPEED MISSION-CRITICAL COMPUTERS WHICH WILL INCLUDE THE VHSIC PROGRAM, AND OUR SOFTWARE INITIATIVES. THESE PROGRAMS ARE COORDINATED CLOSELY WITH DARPA'S STRATEGIC COMPUTING PROGRAM WHICH HAS ALREADY BEEN PRESENTED TO YOU.

LARGE SCALE COMPUTING IN DOD

DOD CURRENTLY HAS LARGE-SCALE COMPUTING MACHINES (SOMETIMES REFERRED TO AS CLASS VI MACHINES) IN THE FOLLOWING INSTALLATIONS:

FACILITY	MACHINE TYPE	APPLICATION
AIR FORCE WEAPONS LAB.	CRAY-1	NUCLEAR WEAPONS EFFECTS AND LASER ANALYSIS
AIR FORCE ARNOLD ENGINEERING DEVELOPMENT CENTER	CRAY-1	WIND TUNNEL DATA REDUCTION, AERODYNAMICS ANALYSIS (THIS MACHINE IS CURRENTLY BEING INTEGRATED WITH THE WIND TUNNEL DATA SYSTEM)
NAVAL FLEET NUMERIC OCEANOGRAPHY CENTER	CYBER-205	OCEANOGRAPHIC AND ATMOSPHERIC RESEARCH
NATIONAL SECURITY AGENCY	CLASSIFIED	COMMUNICATIONS SECURITY, AND FOREIGN INTELLIGENCE

THE ARMY'S BALLISTIC RESEARCH LAB HAS A FOUR PROCESSOR HOMOGENEOUS ELEMENT PROCESSOR SYSTEM BUILT BY DENELCOR CORP. THIS IS A NEW COMMERCIAL MACHINE WHICH IS CURRENTLY CONSIDERED EXPERIMENTAL BY THE USER COMMUNITY. HOWEVER, IT'S PARALLEL ARCHITECTURE IS CONSIDERED TYPICAL OF THE LARGE-SCALE MACHINES OF THE FUTURE. IN ADDITION, SEVERAL DoD LABORATORIES CURRENTLY MAKE USE OF THE LARGE-SCALE COMPUTING COMPLEX AT NASA-AMES RESEARCH CENTER.

THE NAVAL RESEARCH LAB HAS RECENTLY INITIATED PROCUREMENT ACTION FOR ACQUIRING A LARGE-SCALE COMPUTING SYSTEM. THE NAVAL SHIP RESEARCH AND DEVELOPMENT CENTER, THE ARMY'S BALLISTIC MISSILE DEFENSE PROGRAM, THE AIR FORCE GLOBAL WEATHER CENTRAL, AND THE AIR FORCE AERONAUTICAL SYSTEMS DIVISION HAVE PLANS FOR SUCH SYSTEMS IN THE NEAR FUTURE.

DOD'S LARGE-SCALE COMPUTING MACHINES FOR THE MOST PART ARE ENGAGED IN CLASSIFIED ACTIVITIES AND ARE HEAVILY UTILIZED. THEREFORE ACCESS TO THESE MACHINES HAS ESSENTIALLY BEEN LIMITED TO THOSE RESEARCHERS DIRECTLY ENGAGED IN THE WORK OF THE ORGANIZATIONS POSSESSING THEM.

IN ADDITION TO THE APPLICATIONS MENTIONED ABOVE THE DEPARTMENT'S ACTIVITY IN LARGE-SCALE COMPUTING INCLUDES AEROSPACE CALCULATIONS AND ORBITAL MECHANICS, SIGNAL PROCESSING, WEATHER PREDICTION, AND GENERAL MODELING AND SIMULATION. NOW HOWEVER, OTHER APPLICATIONS ARE EMERGING SUCH AS ARTIFICIAL INTELLIGENCE, GRAPHICS, COMPLEX IMAGE PROCESSING, COMMAND AND CONTROL, AND SPACE-BASED DEFENSE.

THE PRESIDENT'S RECENTLY ANNOUNCED INITIATIVE IN SPACE-BASED DEFENSE WILL PERHAPS PRESENT THE MOST SEVERE OF ALL FUTURE MILITARY COMPUTATIONAL REQUIREMENTS. THE REAL TIME COMPUTATIONAL REQUIREMENT IN A VERY HOSTILE ENVIRONMENT FOR DETECTION, TRACKING, AND DISCRIMINATION OF LITERALLY THOUSANDS OF SPACE OBJECTS, INCLUDING DECOYS, ACQUISITION OF THREATENING TARGETS, AND GUIDING DEFENDING WEAPONS TO A VERY LARGE NUMBER OF ENEMY WARHEADS MAY INDEED BE THE MOST FORMIDABLE TECHNOLOGICAL CHALLENGE FACING COMPUTER ENGINEERS AND SCIENTISTS IN THE FORESEEABLE FUTURE. EVEN WHEN WE ARE ABLE IN THE FUTURE TO DESIGN AND PRODUCE COMPUTER ARCHITECTURES WITH THE RAW COMPUTATIONAL POWER REQUIRED, THE DESIGN AND DEVELOPMENT OF THE SOFTWARE TO OPERATE AND CONTROL SUCH SYSTEMS WILL BE FAR BEYOND ANYTHING CONSIDERED PREVIOUSLY. AS ADVANCED AND SOPHISTICATED AS THE APOLLO-SKYLAB AND THE SPACE SHUTTLE ARE IN COMPUTER AND SOFTWARE TECHNOLOGY, THEY WILL BE DWARFED IN COMPARISON.

ANOTHER SIGNIFICANT COMPUTATIONAL REQUIREMENT ASSOCIATED WITH THE DEVELOPMENT OF THE DEPARTMENT'S REAL TIME SYSTEMS IS FOR SIMULATORS TO DRIVE THE OPERATIONAL SYSTEM SOFTWARE. SINCE THESE SIMULATION SYSTEMS MUST BE ABLE TO GENERATE ARTIFICIAL THREATS AGAINST WHICH THE OPERATIONAL SYSTEMS MUST BE TESTED, THEY TOO MUST RUN IN REAL TIME. IT IS NOT UNCOMMON FOR THE SOFTWARE ASSOCIATED WITH SUCH SIMULATORS TO BE AN ORDER OF MAGNITUDE LARGER AND TO CONSUME TEN TIMES AS MUCH COMPUTATIONAL CAPACITY AS THEIR COUNTERPART OPERATIONAL SYSTEMS. CLEARLY, THE SIMULATION REQUIREMENTS FOR OUR FUTURE DEFENSE SYSTEMS, [illegible] AS A STRATEGIC SPACE-BASED DEFENSE SYSTEM, WILL BE TRULY ENORMOUS BY ANY STANDARDS OF MEASURE. EVERY ONE OF OUR MAJOR SYSTEM PROGRAMS MAKES A SIGNIFICANT INVESTMENT IN SIMULATION SOFTWARE, AND OUR BASIC RESEARCH PROGRAMS INCLUDE MUCH EFFORT IN MODELING AND SIMULATION.

THE RECOGNITION OF THE EXTREMELY HIGH RATE OF INCOMING RAW SIGNALS COLLECTED BY MODERN RADAR AND OPTICAL SENSORS, AND THE NEED TO ANALYZE SUCH SIGNALS IN REAL-TIME WAS THE MOTIVATION FOR INITIATING THE VERY HIGH SPEED INTEGRATED CIRCUIT, OR VHSIC PROGRAM IN 1980. A SPACE-BASED DEFENSE SYSTEM IS TYPICAL OF THE HIGH-END SIGNAL PROCESSING REQUIREMENT. SUCH A SYSTEM WOULD LIKELY HAVE MANY THOUSANDS OF OBJECTS IN THE FIELD OF VIEW OF ITS SENSORS AT ONE TIME. THESE SENSORS WOULD HAVE TO HANDLE UNPRECEDENTED AMOUNTS OF DATA AND WILL REQUIRE SPECIALIZED DIGITAL SIGNAL PROCESSING CHIPS THAT ARE EXTREMELY DENSE AND OPERATE AT VERY HIGH SPEEDS. THIS TECHNOLOGY IS NOW BEING DEVELOPED UNDER OUR VHSIC PROGRAM. I WILL DISCUSS SOME OF THE DETAILS OF OUR VHSIC PROGRAM SHORTLY.

ALL OF THE ADVANCEMENTS IN SUPERCOMPUTERS IN THE PAST TWO DECADES HAVE COME FROM INCREASED SPEED IN THE BASIC CIRCUITS, COUPLED WITH GREATER COMPUTATIONAL PARALLELISM. TO ACHIEVE HIGHER SPEED CIRCUITS, PATH LENGTHS HAVE TO BE SHORTENED AND THE DISTANCE OF ELECTRON TRAVEL REDUCED. THIS REQUIRES MUCH HIGHER INTEGRATED CIRCUIT DENSITIES WHICH LEADS TO PROBLEMS IN HEAT REMOVAL. AS CIRCUIT GEOMETRIES ARE REDUCED, INTEGRATED CIRCUITS BECOME MORE DIFFICULT TO PRODUCE AND BECOME MORE SUSCEPTABLE TO FAILURE DUE TO NATURAL RADIATION AND TO IMPURITIES IN

THEIR BASIC MATERIALS. WHILE OTHER MATERIALS, SUCH AS GALLIUM ARSENIDE, PROMISE MUCH GREATER SPEEDS, ALL CIRCUIT TECHNOLOGIES ULTIMATELY FACE FUNDAMENTAL PHYSICAL LIMITATIONS. OUR VHSIC PROGRAM AND DARPA'S VERY LARGE SCALE INTEGRATION OR VLSI RESEARCH PROGRAM DIRECTLY ADDRESS THESE CIRCUIT TECHNOLOGY PROBLEMS.

HIGHLY PARALLEL ARCHITECTURES ULTIMATELY ARE THE BASIS FOR SIGNIFICANTLY GREATER IMPROVEMENTS IN LARGE SCALE COMPUTING IN THE FUTURE. DEVELOPMENT OF SUCH ARCHITECTURES IN THEMSELVES ARE EXTREMELY DIFFICULT UNDERTAKINGS. THIS IS A RELATIVELY NEW AND IMMATURE TECHNOLOGY. MAKING THE TRANSITION FROM USE OF A FEW PROCESSORS TO ACHIEVE PARALLELISM TO A VERY LARGE NUMBER OF PARALLEL PROCESSORS PRESENTS AN EXTREMELY DIFFICULT CHALLENGE. MEETING IT REQUIRES WORK NOT ONLY ON THE PHYSICAL STRUCTURES BUT ALSO ON ALGORITHMS, LANGUAGES, AND SUPPORT SOFTWARE. CONSIDERABLE RESEARCH HAS BEEN CONDUCTED BY THE DEPARTMENT IN DISTRIBUTED DATA PROCESSING ARCHITECTURES WHERE NUMEROUS PROCESSORS OPERATE IN A COOPERATING MANNER AND COMMUNICATE WITH EACH OTHER OVER DATA BUSES. THIS CONCEPT IS EXTREMELY IMPORTANT FOR SURVIVABILITY AND FAIL-SOFT SYSTEM REQUIREMENTS. BUT THERE ARE NO EASY ANSWERS. EFFICIENT USE OF COMPUTING RESOURCES IN SUCH SYSTEMS IS DIFFICULT TO ACHIEVE.

THE ADA PROGRAMMING LANGUAGE WAS DESIGNED FOR EXECUTION ON EITHER SINGLE OR MULTI-PROCESSORS WITHOUT CHANGING THE SOURCE PROGRAM. WHILE THIS IS AN IMPORTANT STEP IN ADDRESSING PARALLELISM, MUCH MORE EFFORT IS NEEDED TO UNDERSTAND HOW TO PARTITION COMPUTING PROBLEMS TO TAKE

ADVANTAGE OF THESE NEW ARCHITECTURES, PARTICULARLY WHERE MANY TENS, HUNDREDS OR EVEN THOUSANDS OF PARALLEL ELEMENTS MAY BE EMPLOYED IN A SINGLE COMPUTATIONAL UNIT.

TOGETHER WITH THE DEPARTMENT OF ENERGY, DOD IS JOINTLY SUPPORTING DEVELOPMENT WORK AT LAWRENCE LIVERMORE NATIONAL LABORATORY ON A HIGHLY PARALLEL SUPERCOMPUTER SYSTEM KNOWN AS THE S-1. THE CURRENT DESIGN OF THE S-1 ALLOWS UP TO SIXTEEN PROCESSORS, EACH OF WHICH IS APPROXIMATELY EQUIVALENT TO A CRAY-1, TO COMMUNICATE WITH SIXTEEN MEMORY MODULES THROUGH A CROSSBAR SWITCH. THIS ARCHITECTURE EMPLOYS WHAT IS CALLED "DANCEHALL" DESIGN WHERE THE PROCESSORS ARE ALIGNED ALONG ONE SIDE AND THE MEMORY MODULES ALONG ANOTHER. A MAJOR PART OF THIS EFFORT HAS BEEN THE DEVELOPMENT OF A POWERFUL COMPUTER AIDED DESIGN CAPABILITY CALLED "STRUCTURED COMPUTER AIDED LOGIC DESIGN" OR SCALD. SCALD IS EXPECTED TO REDUCE SIGNIFICANTLY THE TIME REQUIRED TO DESIGN COMPUTING SYSTEMS. INDUSTRY HAS SHOWN MUCH INTEREST IN SCALD. MORE THAN 160 COPIES HAVE BEEN RELEASED. IT IS BEING USED AS THE BASIS FOR DEVELOPING COMPUTER-AIDED ENGINEERING WORK STATIONS FOR DIGITAL SYSTEMS DESIGN. SEVERAL GENERATIONS OF S-1 PROCESSORS HAVE BEEN DESIGNED USING SCALD. WHILE THE FABRICATION OF THE S-1 SYSTEM IS PROCEEDING SATISFACTORILY, MORE PROGRESS NEEDS TO BE MADE ON APPROACHES TO EXPLOIT THE PARALLELISM OF THIS SYSTEM.

THE DEPARTMENT IS ALSO SUPPORTING RESEARCH IN SYSTOLIC ARRAY PROCESSING WHERE FIXED NETWORKS OF RELATIVELY SIMPLE ARITHMETIC UNITS ARE ARRAYED IN CONFIGURATIONS TO MATCH SPECIFIC PROBLEMS. SUCH SYSTEMS HAVE NO MEMORY REGISTERS OR DATA BUFFERS BETWEEN THE ARITHMETIC UNITS. EACH ARITHMETIC UNIT IN A SEQUENCE PERFORMS A DIFFERENT FIXED SET OF COMPUTATIONS. THESE COMPUTATIONS ARE PERFORMED SYNCHRONOUSLY, IN LOCK-STEP FASHION AND THE RESULTS IMMEDIATELY SHIFTED TO THE NEXT PROCESSOR IN THE PLANNED SEQUENCE. THE TERM "SYSTOLIC" COMES FROM THE CLOCKED OR RHYTHMIC NATURE OF THE PROCESS, ANALOGOUS TO THE HEARTBEAT. WHILE SUCH CONFIGURATIONS WILL BE EXTREMELY FAST FOR SOME APPLICATIONS, TODAY'S STRUCTURES ARE COMPLETELY INFLEXIBLE, REQUIRING AN ENTIRELY NEW DESIGN FOR EACH DIFFERENT APPLICATION.

DOD WILL CONTINUE TO HAVE IMPORTANT APPLICATIONS FOR CONVENTIONAL LARGE SCALE SCIENTIFIC COMPUTING SYSTEMS, AND WILL ALWAYS BE ABLE TO UTILIZE HIGHER PERFORMANCE MACHINES. HOWEVER, IN THE RECENT PAST, WITH THE EXCEPTION OF THE NATIONAL SECURITY AGENCY (NSA) REQUIREMENTS WHICH I'LL DISCUSS SEPARATELY, DOD'S REQUIREMENTS IN THIS AREA HAVE BEEN LESS STRESSING THAN THOSE OF OTHER AGENCIES SUCH AS DOE AND NASA, AND DOD'S NEEDS HAVE GENERALLY BEEN MET BY PRIVATE INDUSTRY DEVELOPMENTS. DOD'S MOST STRESSING NEEDS ARE IN MISSION CRITICAL COMPUTER RESOURCES, A RAPIDLY GROWING AREA THAT I WILL ADDRESS LATER

IN MY PRESENTATION. DOD'S DRIVING REQUIREMENTS IN THE LONG TERM WILL BE DOMINATED BY THE NEED FOR VERY HIGH PERFORMANCE IN BOTH COMPUTATIONAL AND LOGICAL PROCESSING TO OPERATE AND CONTROL EXTREMELY COMPLEX AND SOPHISTICATED REAL-TIME MILITARY SYSTEMS. THESE ARE THE REQUIREMENTS WHICH WILL BE ADDRESSED BY DARPA'S STRATEGIC COMPUTING PROGRAM WHICH HAS ALREADY BEEN COVERED IN THE TESTIMONY OF DR. COOPER, THE DIRECTOR OF DARPA.

BECAUSE OF THE HIGHLY CLASSIFIED NATURE OF NSA'S WORK, IT IS DIFFICULT TO BE VERY SPECIFIC ABOUT THEIR REQUIREMENTS. HOWEVER, THE AGENCY HAS PLAYED A SIGNIFICANT ROLE IN INFLUENCING COMPUTER DEVELOPMENT. IT IS A PRIME USER OF SUPERCOMPUTERS TODAY AND DEPENDS ON SUCH SYSTEMS TO SUCCESSFULLY AND EFFECTIVELY ACCOMPLISH HIGHLY SOPHISTICATED TASKS ASSOCIATED WITH COMMUNICATIONS SECURITY AND FOREIGN INTELLIGENCE. ITS MISSION PERFORMANCE WILL NATURALLY INCREASE AS THE STATE-OF-THE-ART IN COMPUTER TECHNOLOGY EXPANDS. MEMBERS OF THE AGENCY ARE ALREADY WORKING WITH PRIVATE INDUSTRY IN AN EFFORT TO ENHANCE SUPERCOMPUTER TECHNOLOGY.

IN ALL OF THE APPLICATIONS OF COMPUTERS WITH COMPLEX ARCHITECTURES EITHER LARGE OR SMALL WE WILL BE FACING AN EXTREMELY DIFFICULT ISSUE -- THAT OF MATCHING THE PROBLEMS WITH THE MACHINES. WHAT KIND OF ALGORITHMS CAN WE DEVELOP TO EFFICIENTLY UTILIZE THESE ARCHITECTURES? HOW DO WE MAP OUR PROBLEMS ONTO THESE MACHINES, OR HOW DO WE CONFIGURE OUR MACHINES TO ADDRESS THESE PROBLEMS? THERE IS ALSO THE RELATED QUESTION, HOW CAN WE CREATE SOFTWARE TOOLS TO AID OUR EFFORTS IN MANAGING THIS MATCH.

IT IS IMPORTANT THAT A MAJOR PORTION OF EVERY PROGRAM TO DEVELOP A NEW ARCHITECTURE BE AIMED AT ADDRESSING THESE APPLICATION AND SOFTWARE ISSUES. CLEARLY, ADVANCES IN ARCHITECTURES AND HARDWARE, WHETHER FOR LARGE-SCALE OR SMALL-SCALE COMPUTING SYSTEMS, MUST GO HAND-IN-GLOVE WITH COMPARABLE ADVANCES IN OUR SOFTWARE CAPABILITY.

MISSION CRITICAL COMPUTING IN DOD

I WOULD LIKE TO NOW TURN TO THE BROADER BASE OF DOD'S CURRENT COMPUTER REQUIREMENTS. THESE REQUIREMENTS ARISE FROM OUR MORE CONVENTIONAL TACTICAL SYSTEMS. THEY DERIVE FROM THE POLICY THAT WE WILL MEET THE NUMERICAL SUPERIORITY OF THE SOVIET AND EAST BLOC NATIONS WITH TECHNOLOGICALLY SUPERIOR WEAPONS SYSTEMS.

THE REQUIREMENTS ON OUR WEAPONS SYSTEMS ARE SEVERE. THEY MUST STAND OFF FAR FROM THE FRONT LINES YET BE ABLE TO "SEE" DEEPLY INTO ENEMY TERRITORY, HAVE THE CAPABILITY FOR DEEP AND QUICK PENETRATION, HAVE EXTENSIVE RANGE, FLEXIBILITY FOR WIDELY DIVERSE MISSIONS, HIGH RELIABILITY, LONG ENDURANCE, HIGH MOBILITY, AND GREAT LETHALITY. ADVANCES IN TECHNOLOGY ARE CRITICAL TO PROVIDING THESE CAPABILITIES. THE WARSAW PACT NATIONS ARE ALSO ADVANCING IN TECHNICAL SOPHISTICATION. THEY ARE WELL EQUIPPED, AND WELL PREPARED. THERE IS NO TURNING BACK ON OUR EARLY DECISION TO GAIN LEVERAGE FROM OUR SUPERIOR TECHNICAL POTENTIAL. THE OFTEN USED TERM "ARMS RACE" IS A MISNOMER. WE ARE REALLY IMMERSED IN A TECHNOLOGICAL RACE WITH OUR ADVERSARIES. OUR COMPUTER TECHNOLOGY ADVANCES OF THE PAST HAVE GIVEN US CONSIDERABLE LEVERAGE IN THIS RACE BUT WE MUST WORK HARD TO REMAIN AHEAD.

COMPUTERS ARE ESSENTIAL TO PROVIDING THE CAPABILITIES REQUIRED IN OUR ADVANCED SYSTEMS. THE USE OF COMPUTERS [illegible] BOTH TACTICAL AND STRATEGIC MILITARY SYSTEMS IS ALREADY SO WIDESPREAD THAT IT IS DIFFICULT TO ENVISION A FUTURE MAJOR SYSTEM THAT WILL NOT EMPLOY THEM. SOPHISTICATED SENSORS, COUNTERMEASURES, NAVIGATION, GUIDANCE, COMMUNICATIONS, AND COMMAND AND CONTROL, ALL PROVIDED BY MODERN ELECTRONICS, ARE BEING GREATLY ENHANCED BY COMPUTER TECHNOLOGY. SUCCESS ON THE BATTLEFIELD NOW DEPENDS HEAVILY UPON COMPUTER AND SOFTWARE TECHNOLOGY.

OUR FUTURE SYSTEMS DEMAND VERY CAPABLE, HIGHLY COMPLEX, AND EXTREMELY RELIABLE HARDWARE AND SOFTWARE. REQUIREMENTS IN EACH OF THESE AREAS CONTINUE TO GROW RAPIDLY, ALMOST AT AN EXPONENTIAL RATE. CLEARLY, MISSION CRITICAL COMPUTER RESOURCE TECHNOLOGIES ARE ESSENTIAL TO MAINTAINING U.S. WORLD LEADERSHIP IN MODERN DEFENSE WEAPONS AND ARE VITAL TO OUR NATIONAL SECURITY.

I MENTIONED EARLIER THAT THE DEMANDS FOR EXTREMELY HIGH PERFORMANCE SIGNAL PROCESSING CAPABILITY TO REDUCE THE VAST QUANTITIES OF DATA COLLECTED BY MODERN SENSOR SYSTEMS LED TO INITIATING OUR VHSIC PROGRAM. WHILE SPECIALIZED HIGH PERFORMANCE SIGNAL PROCESSING CHIPS WERE THE DRIVING REQUIREMENT, THE VHSIC PROGRAM ALSO IS PROVIDING THE TECHNOLOGY FOR A WIDE VARIETY OF EXTREMELY SMALL, FAST, AND ULTRA-RELIABLE COMPUTER ELEMENTS SUITABLE FOR OPERATION IN SEVERE MILITARY ENVIRONMENTS.

MODERN TACTICAL COMPUTER SYSTEM REQUIREMENTS ARE DOMINATED BY THE NEED FOR SMALL SIZE AND WEIGHT, LOW POWER CONSUMPTION, DURABILITY IN A HOSTILE ENVIRONMENT, HIGH RELIABILITY AND INCREASED PROCESSING THROUGHPUT, PARTICULARLY IN SIGNAL PROCESSING APPLICATIONS. THE VHSIC PROGRAM WAS INITIATED TO ACHIEVE A HUNDREDFOLD IMPROVEMENT IN MOST OF THE ABOVE AREAS. THE PROGRAM IS PROCEEDING THROUGH TWO PHASES OF TECHNOLOGY DEVELOPMENT. PHASE 1 NOW APPROACHING COMPLETION IS ACHIEVING 1.0 TO 1.25 MICRON GEOMETRIES, RESULTING IN THE ABILITY TO PACK 20,000 TO 50,000 LOGIC GATES OR THE EQUIVALENT OF ABOUT 150,000 TRANSISTORS ON A SINGLE 1/4 INCH BY 1/4 INCH (APPROX.) CHIP. PHASE 2 WILL EXTEND THIS TECHNOLOGY TO SUBMICRON GEOMETRIES, ALLOWING THE PACKING OF UP TO 100,000 LOGICAL GATES ON THE SAME SIZE CHIP. THIS PROGRAM IS DEVELOPING THE TECHNOLOGY TO PRODUCE CUSTOM MICROELECTRONIC CHIPS FOR UNIQUE APPLICATIONS OF MILITARY SYSTEM.

THERE ARE STILL OBSTACLES TO BE OVERCOME IN THE COST EFFECTIVE PRODUCTION OF VHSIC TECHNOLOGY, BUT NONE ARE INSURMOUNTABLE. TODAY, IT APPEARS OUR MAJOR BOTTLENECK TO WIDESPREAD USE OF THIS TECHNOLOGY IS IN THE DESIGN OF NEW APPLICATION-UNIQUE CHIPS. OUR EFFORTS BOTH IN IMPROVING THE MANUFACTURING YIELD AND IN COMPUTER AIDED DESIGN POSE SIGNIFICANT CHALLENGES AND OFFER TREMENDOUS OPPORTUNITIES. OUR FREE WORLD LEADERSHIP IN SEMICONDUCTORS IS SLIM. IMPROVED CAPABILITIES IN THESE AREAS ARE VITAL IF WE ARE TO ASSURE U.S. SUPREMACY IN THIS CRITICAL TECHNOLOGY.

THE GROWTH IN THE USE OF DIGITAL COMPUTERS IN WEAPONS SYSTEMS STEMS FROM TWO BASIC FACTORS. FIRST, THE INHERENTLY BETTER COMPUTATIONAL PERFORMANCE OF DIGITAL DEVICES OVER THEIR COUNTERPART ANALOG DEVICES, AND SECOND, THE FLEXIBILITY PROVIDED BY SOFTWARE. WE HAVE CLEAR EVIDENCE THAT EVEN WITH THE PROBLEMS WE HAVE ENCOUNTERED WITH SOFTWARE, THE COST TO MODIFY A WEAPON SYSTEM FUNCTION OR PERFORMANCE VIA A SOFTWARE CHANGE IS LIKELY TO BE SEVERAL ORDERS OF MAGNITUDE LESS THAN IF THE HARDWARE MUST BE MODIFIED.

EXPLOITATION OF THE FUNDAMENTAL BENEFITS OF DIGITAL COMPUTERS AND SOFTWARE HAS BEEN ACCOMPANIED BY A NEW SET OF PROBLEMS. THESE INCLUDE DIFFICULTIES IN PROVIDING LOGISTICS SUPPORT FOR THE PROLIFERATION OF A WIDE VARIETY OF COMPUTER TYPES, AND INCREASING SOFTWARE COST, SCHEDULE, AND PERFORMANCE PROBLEMS CAUSED BY THE RAPIDLY INCREASING DEMAND FOR NEW, LARGER, AND MORE COMPLEX SOFTWARE IN THE FACE OF MUCH SLOWER GROWTH IN OUR CAPABILITY TO DEVELOP IT. SINCE THE GROWTH IN SOFTWARE INVESTMENT IN DEFENSE SYSTEMS IS OFTEN MISUNDERSTOOD AS INCREASING COST FOR THE SAME LEVEL OF SOFTWARE EFFORT, I THINK IT IS IMPORTANT TO SET THE RECORD STRAIGHT AND EXPLAIN FOR A MINUTE HOW THIS GROWTH IS OCCURRING. THERE ARE FOUR DISTINCT FACTORS DRIVING THIS GROWTH.

FIRST, WE ARE RAPIDLY MIGRATING FUNCTIONS TO SOFTWARE WHICH WERE PREVIOUSLY PERFORMED WITH HARD-WIRED DEVICES, FOR EXAMPLE, IN COMMUNICATION SWITCHING.

SECOND, ENTIRELY NEW FUNCTIONS, IMPOSSIBLE BEFORE THE EMERGENCE OF STORED PROGRAM DIGITAL DEVICES, ARE NOW BEING DEVELOPED. AN EXAMPLE OF THIS IS THE DIGITAL MAP DISPLAY NOW UNDER DEVELOPMENT FOR AIRCRAFT.

THIRD, TO REDUCE WEIGHT, SPACE, POWER CONSUMPTION, COST, ET CETERA, FUNCTIONS PREVIOUSLY PERFORMED IN SEPARATE SUBSYSTEMS ARE BEING COMBINED INTO SINGLE SUBSYSTEMS USING TIME-SHARED COMPONENTS. THE AIR FORCE'S INTEGRATED COMMUNICATIONS, NAVIGATION, IDENTIFICATION AVIONICS PROGRAM, OR ICNIA, IS AN EXCELLENT EXAMPLE.

FOURTH, THE GROWING USE OF SEPARATE, DIGITALLY IMPLEMENTED FUNCTIONS OPERATING IN PARALLEL, AND COMMUNICATING VIA MULTIPLEX BUSES, IMPOSES NEW CONTROL REQUIREMENTS WHICH MUST BE IMPLEMENTED IN SOFTWARE FORM.

IN 1975 WE TOOK AGGRESSIVE ACTION TO DEAL WITH THE WORSENING SOFTWARE SITUATION VIA OUR ADA INITIATIVE. MOTIVATED BY THE EXCESSIVE USE OF ASSEMBLY LANGUAGE, THE PROLIFERATION OF HIGH ORDER LANGUAGES, EACH OF WHICH WAS POORLY SUPPORTED, THE LACK OF LANGUAGES ORIENTED TO MILITARY REAL-TIME NEEDS, AND EXCESSIVE MACHINE-DEPENDENCE, THE DEPARTMENT OF DEFENSE SET OUT TO ACHIEVE A COMMON HIGH ORDER LANGUAGE FOR MISSION-CRITICAL SYSTEMS, I.E., THOSE SYSTEMS USED FOR PURPOSES OTHER THAN ADP, OR ADMINISTRATIVE COMPUTING.

EVOLVED FROM AN ITERATIVELY REFINED SET OF COMMON REQUIREMENTS, ADA TODAY CONTAINS MANY FEATURES THAT PROVIDE POWERFUL SUPPORT FOR THE DEVELOPMENT AND MAINTENANCE OF LARGE-SCALE, REAL-TIME SOFTWARE SYSTEMS. THE POTENTIAL BENEFITS DUE TO THE USE OF ADA HAVE BEEN ESTIMATED, VIA THREE INDEPENDENT STUDIES, TO BE IN EXCESS OF $100 MILLION PER YEAR.

ADA WAS DEVELOPED AND MATURED WITH EXTENSIVE SUPPORT FROM THE COMMUNITY AT LARGE. FIRST, A MILITARY STANDARD, IT THEN EVOLVED INTO A JOINT ANSI AND MIL STANDARD WHICH WAS PUBLISHED THIS PAST FEBRUARY. THERE ARE ALREADY OVER 25 ADA COMPILER EFFORTS IN THE U.S., EUROPE, AND JAPAN. THREE OF THESE WERE RECENTLY VALIDATED AND CERTIFIED BY OUR ADA JOINT PROGRAM OFFICE. WE ARE VERY PROUD OF THE PROGRESS THAT HAS BEEN MADE. WE ARE ALSO VERY PLEASED WITH THE WIDESPREAD ACCEPTANCE AND THE VERY STRONG SUPPORT ADA NOW HAS FROM THE CONGRESS, THE UNITED STATES INDUSTRY, THE EUROPEAN ECONOMIC COMMUNITY, NATO, THE PROFESSIONAL SOCIETIES, AND ACADEMIA. TO DATE, OVER 20 TEXTBOOKS ON ADA HAVE BEEN PUBLISHED AND HUNDREDS OF ARTICLES AND PAPERS HAVE APPEARED. ALSO, MANY COURSES ARE NOW BEING OFFERED ON ADA AT BOTH GRADUATE AND UNDERGRADUATE LEVELS.

THE MATURITY OF THE LANGUAGE AND THE EMERGING COMPILERS HAVE MOTIVATED US TO ISSUE A POLICY MAKING ADA MANDATORY FOR FUTURE MISSION-CRITICAL APPLICATIONS. THIS POLICY REQUIRES THAT ADA BE USED IN ALL SYSTEMS ENTERING THE ADVANCED DEVELOPMENT PHASE AFTER JANUARY 1, 1984 AND IN ALL SYSTEMS ENTERING THE FULL-SCALE ENGINEERING DEVELOPMENT PHASE AFTER JULY 1, 1984. WE ALREADY HAVE OVER 15 SYSTEMS LINED UP TO USE ADA.

AN ADA-RELATED INITIATIVE THAT IS VERY PROMISING IS THE DEVELOPMENT OF INTERFACE STANDARDS FOR ADA PROGRAMMING SUPPORT ENVIRONMENTS. THESE WILL ALLOW USE OF A RANGE OF ADA COMPILERS BY FACILITATING THE MOVEMENT OF SOFTWARE TOOLS AND DATA ACROSS ENVIRONMENTS WITHOUT CONVERSION OR REPROGRAMMING.

ADA HAS ARRIVED. BUT DOD DID NOT DO IT ALONE. WITHOUT THE HELP WE RECEIVED FROM THE ACADEMIC AND INDUSTRIAL COMMUNITIES, WE PROBABLY WOULD NOT HAVE SUCCEEDED AT ALL.

WHILE THE BENEFITS FROM THE USE OF ADA, BOTH IN INITIAL DEVELOPMENT AND IN LATER REUSE AND TRANSPORTABILITY OF SOFTWARE, ARE OBVIOUS; WE REALIZE THAT ADA IS NOT A PANACEA. WHILE ITS STRENGTH AS A TECHNICAL TOOL PRIMARILY FOR USE IN THE IMPLEMENTATION (CODING) PHASE OF SOFTWARE DEVELOPMENT IS UNQUESTIONED, WE NOW REALIZE THAT IMPROVEMENTS ARE ALSO BADLY NEEDED IN OTHER PHASES OF THE SOFTWARE LIFE CYCLE AND IN ACTIVITIES OTHER THAN IMPLEMENTATION.

ABOUT A YEAR AND A HALF AGO, AFTER AN IN-DEPTH STUDY OF THE SOFTWARE PROBLEM, WE CONCLUDED THAT A BROADER ATTACK WAS NEEDED IF WE WERE TO CONTROL THE SOFTWARE "EXPLOSION." WE CALLED THIS ATTACK THE SOFTWARE INITIATIVE AND LATER ADDED THE NAME STARS FOR "SOFTWARE TECHNOLOGY FOR ADAPTABLE, RELIABLE SYSTEMS." THE INITIAL STRATEGY WAS COMPLETED IN OCTOBER 1982 AND WAS THEN REFINED BY A JOINT TASK FORCE BY MARCH 1983.

IT HAS BEEN SUGGESTED IN THE PRIVATE SECTOR THAT THE STARS PROGRAM WILL MEAN TO DOD AND THE U.S. SOFTWARE INDUSTRY WHAT VHSIC HAS AND WILL MEAN TO DOD AND THE U.S. SEMICONDUCTOR INDUSTRY. VHSIC AND STARS ARE COMPLEMENTARY PROGRAMS. BOTH ARE ESSENTIAL TO MEETING THE OBJECTIVES OF THE PRESIDENT'S NATIONAL SECURITY DECISION DIRECTIVE-85 (NSDD-85).

ALTHOUGH THE "SOFTWARE PROBLEM" HAS BEEN APPARENT SINCE THE EARLY 1970s, IT HAS ONLY BEEN THROUGH THE ADA HIGH ORDER LANGUAGE PROGRAM, AND MORE RECENTLY THE Ada PROGRAMMING SUPPORT ENVIRONMENTS, THAT WE BEGUN TO SEE A POSITIVE EFFECT ON THIS INSIDIOUS AND PERVASIVE STUMBLING BLOCK. A MORE GLOBAL REVIEW OF OUR REQUIREMENTS IN THE MISSION CRITICAL COMPUTER RESOURCE AREA HAS REVEALED THAT THERE IS NO SINGLE OR SIMPLE SOLUTION. PROGRESS WILL BE ACHIEVED ONLY BY MAKING PROGRESS ON A NUMBER OF FRONTS SIMULTANEOUSLY.

IN THE MOST GENERAL SENSE, OUR GOAL IS TO IMPROVE PRODUCTIVITY WHILE ACHIEVING GREATER SYSTEM RELIABILITY AND ADAPTABILITY. THAT IS, WE WANT TO MAKE DEVELOPMENT AND POST DEPLOYMENT SUPPORT FASTER, LESS EXPENSIVE AND MORE PREDICTABLE. THIS IS INTENDED TO RESULT IN SOFTWARE THAT IS MORE POWERFUL, RELIABLE AND ADAPTABLE. OUR TASK IS TO CREATE THE TECHNOLOGY AND ESTABLISH THE PRACTICES TO IMPROVE MISSION-CRITICAL DEFENSE SOFTWARE AND PRESERVE THE U.S. LEAD.

THE STARS PROGRAM FOCUSES ON CREATION OF AN INTEGRATED, HIGHLY AUTOMATED DEVELOPMENT AND SUPPORT ENVIRONMENT FOR SOFTWARE. SUCH AN ENVIRONMENT IS TERMED A SOFTWARE ENGINEERING ENVIRONMENT AND IS ANALOGOUS TO THE SET OF TOOLS AND CAPABILITIES KNOWN AS CAD/CAM FOR COMPUTER AIDED DESIGN AND COMPUTER AIDED MANUFACTURING FOR VLSI CHIP DEVELOPMENT AND PRODUCTION. WE CANNOT EXPECT TO ACHIEVE A SINGLE SET OF AUTOMATED TOOLS FOR APPLICATION IN ALL SYSTEM AREAS. THEREFORE, EXPLORATION AND EXPERIMENTATION WITH A NUMBER OF SUCH SETS IS PLANNED IN THE PROGRAM. EACH SET MIGHT BE CONSIDERED TO EMPLOY A PARTICULAR METHODOLOGY. WHILE DIFFERENT METHODOLOGIES AND DIFFERENT TOOL SETS WILL BE SUPPORTED, A COMMON FRAMEWORK AND COMMON ARCHITECTURE ARE DEEMED NECESSARY. IMPORTANT IN THE STARS PROGRAM IS THE EFFORT TO FORMALLY OR EXPLICITLY DEFINE A SOFTWARE ENGINEERING ENVIRONMENT. A TRI-SERVICE TEAM IS NOW PREPARING PRELIMINARY REQUIREMENTS FOR SUCH AN ENVIRONMENT. A CRITICAL ELEMENT OF THIS ACTIVITY IS THE ESTABLISHMENT OF STANDARD SOFTWARE INTERFACES TO PROVIDE FOR THE INTEGRATION OF MANY AND VARIED TOOLS NEEDED IN SUCH A SYSTEM.

THE ADA PROGRAM IS THE CORNERSTONE OF THE STARS PROGRAM. ADA IS THE HIGH ORDER PROGRAMMING LANGUAGE WHICH WILL BE EMPLOYED. AT FIRST, AN EXTENDABLE CORE SOFTWARE ENGINEERING ENVIRONMENT WILL BE DEVELOPED BY CAPITALIZING ON THE ON-GOING ADA PROGRAMMING SUPPORT ENVIRONMENT (APSE) EFFORT. THIS CORE ENVIRONMENT WILL BE USED AS A BASIS FOR CONSOLIDATING EXISTING TOOLS THAT HAVE ALREADY BEEN DEVELOPED AND DEMONSTRATED.

THE RELATIVE FOCUS OF ADA AND STARS WITH RESPECT TO THE SOFTWARE LIFE CYCLE IS ILLUSTRATED IN THE CHART BELOW.

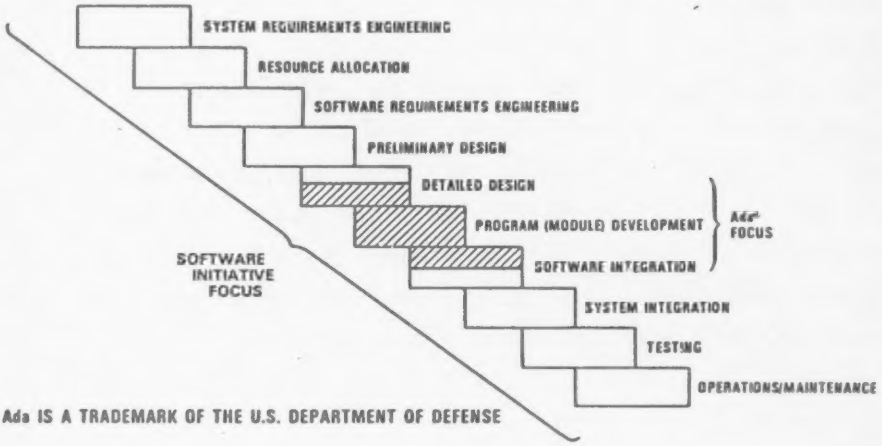
SYSTEM REQUIREMENTS ENGINEERING
RESOURCE ALLOCATION
SOFTWARE REQUIREMENTS ENGINEERING
PRELIMINARY DESIGN
DETAILED DESIGN
PROGRAM (MODULE) DEVELOPMENT
SOFTWARE INTEGRATION
SYSTEM INTEGRATION
TESTING
OPERATIONS/MAINTENANCE
Ada* FOCUS
SOFTWARE INITIATIVE FOCUS

* Ada IS A TRADEMARK OF THE U.S. DEPARTMENT OF DEFENSE

THE STARS PROGRAM AIMS AT DEVELOPING AND INTEGRATING TOOLS TO BE USED IN EARLIER PHASES OF THE LIFE-CYCLE. EXISTING DATA SHOWS THAT CORRECTION OF SPECIFICATION AND DESIGN ERRORS ARE FAR MORE COSTLY WHEN FOUND AFTER SYSTEM EMPLOYMENT, THAN ARE CODING ERRORS. THEREFORE, THERE IS AN EMPHASIS IN STARS TO ESTABLISH IN THE ENVIRONMENT TOOLS WHICH AID IN THE REQUIREMENTS DEFINITION, AND SPECIFICATION PHASES OF THE SOFTWARE LIFE-CYCLE AS SHOWN ON THE CHART ABOVE. THIS EMPHASIS REFLECTS A PHILOSOPHY OF ERROR AVOIDANCE IN THE SOFTWARE LIFE CYCLE OVER ONE OF ERROR CORRECTION. SINCE MOST OF WHAT IS INAPPROPRIATELY CALLED "SOFTWARE MAINTENANCE" IS REALLY ADDITIONAL DEVELOPMENT TO MEET THE NEED FOR UPGRADES AND ENHANCEMENTS -- THIS EMPHASIS IS PROPERLY PLACED. RELATIVELY LITTLE CAN BE DONE TO EASE THE DIFFICULTY IN SUPPORTING THE LARGE INVENTORY OF EXISTING, OFTEN "OLD" SOFTWARE UNTIL IT IS REDEVELOPED USING MODERN SOFTWARE TOOLS AND PRACTICES.

AFTER THE INITIAL VERSIONS HAVE BEEN DEVELOPED AND APPLIED, THE POWER AND SCOPE OF THE AUTOMATED ENVIRONMENT WILL BE FURTHER ENHANCED BY INTRODUCING AND INTEGRATING TOOLS THAT SUPPORT PERSONNEL OTHER THAN THE DESIGNERS AND DEVELOPERS (E.G. PROGRAM MANAGERS). INCLUDED WILL BE THE CAPABILITY TO AUTOMATICALLY COLLECT DATA TO AID IN MEASUREMENT AND EVALUATION OF TOOLS PERFORMANCE.

ADVANCEMENT IN SOFTWARE TECHNOLOGY HAS BEEN OCCURING VIA OUR TECHNOLOGY BASE PROGRAMS, VIA UNIVERSITY RESEARCH CENTERS, AND VIA

INDUSTRY IR&D. HOWEVER, BRINGING THESE ADVANCEMENTS INTO PRACTICE HAS BEEN A VERY SLOW AND LABORIOUS PROCESS. INDIVIDUAL PROGRAMS GENERALLY DO NOT HAVE ADEQUATE RESOURCES TO FULLY EVALUATE RESEARCH 'PRODUCTS' AND CONVERT THEM TO PRODUCTION QUALITY TOOLS. AS PART OF THE SOFTWARE INITIATIVE -- STARS PROGRAM, WE PLAN TO ESTABLISH A SOFTWARE ENGINEERING INSTITUTE TO FOCUS ON THE PROBLEM OF TRANSITIONING NEW SOFTWARE RESEARCH RESULTS INTO PRACTICE IN DOD SYSTEMS. WE NEED TO MAKE PRODUCTS OF OUR RESEARCH EFFORTS AND TO PUT THESE PRODUCTS TO USE. THE RESULTS WILL BE IMPORTANT TO ALL AREAS OF COMPUTING LARGE-SCALE, MISSION-CRITICAL, AND EVEN ADMINISTRATIVE DATA PROCESSING.

SUMMARY

IN SUMMARY, DOD HAS HAD A LONG AND DEEP INVOLVEMENT IN NEARLY ALL ASPECTS OF COMPUTING TECHNOLOGY. FOR THE MOST PART INDUSTRY HAS BEEN MEETING OUR NEEDS FOR CONVENTIONAL LARGE COMPUTING SYSTEMS. WHILE VERY LARGE SCALE NUMERIC COMPUTING MACHINES HAVE BEEN VERY IMPORTANT TO DOD, IN RECENT YEARS THEY HAVE NOT BEEN THE PRINCIPAL FOCUS OF DOD'S DEVELOPMENTS IN COMPUTING. OUR NEEDS ARE CHANGING, AND TECHNOLOGY IS ADVANCING. THE PROMISING GAINS IN SOME AREAS OF ARTIFICIAL INTELLIGENCE ARE PRESENTING NEW OPPORTUNITIES FOR FAR MORE SOPHISTICATED SYSTEMS THAN ARE ACHIEVABLE TODAY. DARPA'S PROGRAM IS AIMED AT THESE OPPORTUNITIES.

DOD'S DRIVING REQUIREMENTS IN COMPUTING ARE DERIVED FROM THE NEEDS FOR MISSION CRITICAL SYSTEMS. THE VHSIC PROGRAM HAS ADDRESSED THE OBJECTIVE OF SHRINKING THE SIZE, WEIGHT AND COST FACTORS TO MEET MILITARY UNIQUE ENVIRONMENTS AND APPLICATIONS. OUR SOFTWARE INITIATIVES, ADA AND STARS, HAVE BEEN ESTABLISHED TO ADDRESS OUR GROWING SOFTWARE DEMANDS. CLEARLY, THE PRODUCTS OF THIS INITIATIVE, FOR EXAMPLE, AUTOMATED SOFTWARE TOOL SYSTEMS HAVE APPLICABILITY TO BOTH LARGE-SCALE COMPUTING AND MISSION CRITICAL COMPUTING.

UNDENIABLY, IMPROVEMENT IS BEING MADE IN SOME SOFTWARE RELATED AREAS. HOWEVER, ADVANCEMENT IN COMPUTER HARDWARE ARCHITECTURES BRINGS WITH IT THE NEED FOR COUNTERPART ADVANCEMENT IN SOFTWARE TECHNOLOGY. THE TWO ARE NOT SEPARABLE, IF THE HARDWARE ADVANCES ARE TO BE EXPLOITED. IN PARTICULAR MUCH RESEARCH NEEDS TO BE DONE TO UNDERSTAND HOW TO MATCH NON-CONVENTIONAL ARCHITECTURES, PARTICULARLY THE EMERGING HIGHLY PARALLEL ARCHITECTURES TO CONVENTIONAL COMPUTING PROGRAMS, AND HOW TO DEVELOP SOFTWARE TOOLS TO FACILITATE THIS MATCHING. WE WILL BE FOCUSING ON THESE PROBLEMS IN OUR FUTURE RESEARCH PROGRAMS.

DOD'S MISSION-CRITICAL COMPUTING REQUIREMENTS ARE SOMEWHAT UNIQUE. WE MUST PARTITION OUR LARGE-SCALE COMPUTING REQUIREMENTS IN THESE SYSTEMS, AND DISTRIBUTE THE FUNCTIONS AMONG A NUMBER OF PROCESSORS IN ORDER TO ACHIEVE A FAIL-SOFT EFFECT FOR SURVIVABILITY. WE ARE CONTINUING TO WATCH VERY CLOSELY THE DEVELOPMENTS IN

SUPERCOMPUTERS IN THE COMMERCIAL WORLD. THEIR SHRINKING SIZE CERTAINLY DEMANDS THEY BE CONSIDERED FOR DIRECT APPLICATION IN SOME MISSION-CRITICAL SYSTEMS. IF THEY CAN MEET THE OPERATIONAL ENVIRONMENT REQUIREMENTS, AND DEMONSTRATE THE ROBUSTNESS NEEDED FOR SURVIVABILITY, WE BELIEVE SUCH APPLICATIONS MAY SOON BE FEASIBLE.

FINALLY, I WOULD LIKE TO SAY IT IS CRITICAL THAT THE U.S. MAINTAIN ITS LEAD IN SUPERCOMPUTING TECHNOLOGY. DOD BENEFITS AND WILL CONTINUE TO BENEFIT GREATLY FROM PROGRESS IN THIS AREA. WHILE OUR PROGRAMS HAVE NOT BEEN DIRECTLY FOCUSED ON SUPERCOMPUTERS, WE BELIEVE THEY HAVE BEEN SUPPORTIVE. AS OUR REQUIREMENTS CHANGE AND AS TECHNOLOGY ADVANCEMENTS PRESENT NEW OPPORTUNITIES, WE CONTINUE TO ADJUST OUR PRIORITIES AND OUR OBJECTIVES. OUR INITIATIVES CLEARLY DEMONSTRATE THAT.

THANK YOU MR. CHAIRMAN.

Mr. **Fuqua**. Thank you very much, Dr. Martin.

Dr. Knapp, NSF has taken a position, as I understand it, of treating supercomputer as kind of a scarce resource in itself. Is that rather accurate?

Dr. **Knapp**. I think that's correct, yes. It's—I think to characterize it slightly differently, it's taken a position that a supercomputer is like any other large instrument which is necessary for scientific work and that the time on that instrument should be allocated by the peer review process and by—and by—and that we should support such instruments as a part of our overall function of the foundation for supporting academic science.

Mr. **Fuqua**. Well, yesterday we had testimony from Dr. Kozmetsky, which forecast several thousands of units in the next 10 years or so. Dr. Wilson from Cornell testified last—I guess earlier this year, a similar-type thing, that there would be a great demand for it. Dr. Beering testified yesterday, the president of Purdue, that they've had their machine about a year, less than a year, and already had 200 users on the machine and anticipated to go to 5,000 within—I don't remember the timeframe, another year or two, make it accessible to more and more students.

It seems like that that is a significant conflict in where we are going and from what the position that NSF was taking and what other witnesses have alluded to.

Dr. **Knapp**. I think that the NSF is moving as rapidly as possible to get supercomputer time for academic users in 1984 and to try to establish the real need for supercomputers on the basis of workshops and meetings and advisory committees to us.

I think that a number of 1,000 supercomputers in 10 years is perhaps a little too much. I think that we don't know what that number is. I know that the supercomputers will have a large number of users. Five thousand for the Purdue system may be realistic. It may not be. I think we are trying very hard to establish that base line, and it's very difficult to establish.

As I said in my testimony, there are areas of science where the people don't know what it is they are not doing because they've never had any access to these machines so they can't really judge what the power of these machines could bring to them.

My own personal feeling is that it—that the need is really quite large. However, I think a reasonably prudent approach to understanding what the need is and to making plans for the future is to go in the sequence that we propose to go through; that is, to get the input from the community itself and to try to understand what we can do to satisfy their needs.

Of course, the other—the other side of that is that we have limited resources in the foundation and we have to understand what we support and what we don't support if we support more supercomputers, and that—this balancing is what we are trying to do with our studies right now.

Mr. **Fuqua**. Well, even if the projections of 5,000 were only half right, that is still quite a market.

Dr. **Knapp**. That is quite a market.

Mr. **Fuqua**. And I recall in I guess the early fifties or midfifties when the Interstate Highway System was proposed and considered by the Congress, when you look today how off they were on the projected traffic that would be utilizing those interstate highways, we may be in a similar-type situation with supercomputers. It is sometimes—someone referred it is like when you start cutting trees in a forest, the more you cut the further you can see, the more uses that are developed for supercomputers.

Dr. **Knapp**. We would like to evolve an organization and a system within the Foundation that can expand as we determine this need to fill whatever needs we feel the academic community requires, both in terms of how we go about funding the operation of centers and how we go about assigning time to university users.

Mr. **Fuqua**. Before I forget, without objection Dr. Martin's statement will be made part of the record in its entirety and, of course, in addition to the summary.

Dr. Martin, what is the status of the university preparedness to the innovation in software and hardware architecture you referred to in your statement?

Dr. **Martin**. You mean the capability of our universities to do work that is supportive of advancements in the areas that I have described?

Mr. **Fuqua**. Correct.

Dr. **Martin**. Well, hopefully it's very good because we are going to rely on them. But I definitely feel that there is an awful lot of good work that has been done in the universities which is in a research product form and it is not in a usable form at the present time. So, there are some things that are sitting on shelves in universities which could be used to help.

Underwriting additional work in the universities would help us in a number of ways. One, I think that they're quite capable of doing the research work necessary to help solve these problems and two, by giving them the opportunity to solve these problems they will find new frontiers for us and move us forward in other dimensions.

Mr. **Fuqua**. How about the training that they will provide for people like your department and others in DOD, the Government, as well as industry?

Dr. **Martin**. Well, I will not hesitate to say that we have some very severe problems in our universities today, only part of which are affecting the computer area. But equipment is one, number of faculty is another. You know, we are short at least 200 faculty in our computer sciences across the United States. Now, that is a very large number. So, having just the faculty to support the education process is a problem.

Mr. **Fuqua**. But where does the faculty get its education? How does the faculty become prepared?

Dr. **Martin**. Within the very universities that we have, but the rate at which we can do this is constrained by the—by the size of the culture that we start with. And we haven't done anything to make a major step change in the size of that culture. And unless we do, we are going to continue to have a deficiency in the number of educators and the number of people who can be effectively trained.

Right now the student-to-professor ratio in the universities is something of which we ought to be appalled. Now, NSF I am sure can give some explicit statistics on that, but I know in my own home university of Georgia Institute of Technology the balance is way out of proportion and we are having to turn away students because we either in that instance educate a lot of students and sacrifice quality or have quality and fewer students.

We've got—you know, we've got some problems in the computer area that have to be addressed in a number of different directions. One is in increasing the faculty, and part of the problem you have in doing that is salaries, not so much positions but to some extent positions, equipment that is state-of-the-art equipment that allows them to do research and to educate their students on modern systems, and it's frustrating to them not to be able to do that and it does turn people away from the faculty situation. So, you know, there are some problems that need to be attacked nationwide in that area.

You know, we don't in most schools have adequate facilities for training students. We certainly don't provide enough money for the kinds of hands on experience that students need with a large variety of equipment that they should have in order to move the technology forward.

Now, if I might, I would like to digress a little bit to an answer that was being given on exposure to supercomputers in the universities. Universities are always equipment constrained, and when they made their decisions as supercomputers were coming along let's say about 10 years ago, usually you had a central computing facility on campus.

That central computing facility was selected by some team that comprised the researchers, the administrative group and the educators, and the systems that were acquired by the universities were a compromise. They weren't terribly good for anybody.

They were time shared systems that could be used by the instructors to let their students get on the machine and have some access to the equipment. They had to satisfy the administrative requirements, and so it took away from the speed of response to the students who were on the machine.

And then there were the people who had research activities and need to do large computations. What they found was that their attention time on that machine was so fragmented it could take them a week to run a computation that had they been given that whole machine in raw form could have been done in a half hour.

And so the situation on the campuses is one that derives from equipment constraints. And as far as how many users are out there, I would say that if you give them the tool, they will use it, and if you give them the tool to use, they will find new uses for it.

But there are reasons why we haven't been able to do that in the past. Part of it is the university funding for machines. Part of it has been the tremendous expense of those machines.

You know, economies of scale do work in the computer area. I remember paying $460 for a scientific hand-held calculator when they first came out. I remember paying $85 for a five-function hand-held calculator, and today you can have them for $4 or $5. So, economies of scale do apply and to what we currently call supercomputers, they will apply there as well.

There are many areas in which supercomputers should be prolific on our campuses. Our physics departments should have a supercomputer all their own. The costs are coming down. Now, that is one dimension that is going to help our universities, but it is only going to help our universities if in fact we give them some money to buy that equipment and maintain it. You know, we can find people who will give them the equipment, but they can't afford the 1 percent it would cost to maintain the equipment.

The other problem that is starting to be resolved is that that equipment came, as I mentioned in my summary, it came without any support software, so you had to have the buy-in price, and that was the time and expertise to develop the software so that you could use the machine. And there aren't an awful lot of Ph.D. students who want to stand around building a bridge so they can walk over it. And we have had to do it out of necessity, but others find other alternatives.

The supercomputer manufacturers today are starting to develop support software for their systems, and what that means is that people who have problems that can be solved by supercomputers can now use those machines to solve their technical or research problems without becoming a computer scientist first and building a tool that's necessary to get where they ultimately want to go.

So, that is turning around. It could use some help, and if there is aid in that area, we will expedite the speed with which supercomputers are available to a larger user base. Again, economies of scale will help because the more people who use these machines and have them available the more people who will generate ideas for how to make that machine friendly.

Mr. **Fuqua**. We addressed that in legislation that came out of this committee and passed the House in March of this year. It's now still languishing in the other body, but we hope to get it moving and address some of the things you are talking about.

The point I was trying to make was that universities are the training ground. That's the only place we train scientists and researchers. They take that knowledge and go other places, to national labs or industry, and they utilize that.

Dr. **Martin**. That's right.

Mr. **Fuqua**. They also make—become researchers and teachers at a university. So, the point I am making, we can't increase the faculty unless we have more facilities and better attraction for people to come into that.

Dr. Knapp, what's probably the most effective approach that NSF can take toward making supercomputers available to university researchers?

Dr. **Knapp**. I think the approach that we have proposed is the most effective approach; that is, to first very quickly buy supercomputer time and allocate it to the scientific disciplines on the basis of peer review, and then on the longer haul to develop any university's computer centers with network to the entire scientific university system of the country.

Mr. **Fuqua**. Where would the computers be located? At certain designated universities?

Dr. **Knapp**. Our plans are now to locate them at designated universities, each center at a university being sized so that it's optimum for student training and also for the kind of consulting and advice that's necessary to the users in order to have optimum use of the facilities.

Mr. Fuqua. And they use networking type——

Dr. Knapp. And then use networking to have access to the entire communities of the scientific disciplines. I think that will work quite well. It's also very expandable. It can be built to a certain size and can be continued larger if the demand is there and the resources are there to do so.

Mr. Boehlert. Would the chairman yield?

Mr. Fuqua. Yes.

Mr. Boehlert. Dr. Knapp, I'd like to ask you, how are we coming with the study on networking? I think that's due before this committee by January of 1984. Are you proceeding with that on course?

Dr. Knapp. Well, as I reported in my testimony, we are starting the purchase of computer time for this year, so we are—yes, there is a study going forward, and you will have a report to the committee, but we are also at the same time starting a program within the foundation for access to supercomputers by university researchers.

Mr. Boehlert. But that report that will come to the committee, will that address the question of networking for the university research centers?

Dr. Knapp. Yes.

Mr. Boehlert. Thank you.

Mr. Fuqua. Mr. Winn.

Mr. Winn. Thank you, Mr. Chairman.

Dr. Knapp, we are all aware of the three operating university supercomputer centers in the United States. What did the acquisition of these centers have in common besides requiring no Federal funds? Did they just go their own way?

Dr. Knapp. I'm not completely familiar with the details of the acquisition, but I think in this case they went their own way. They got—they got through a very fine entrepreneurial spirit; these people developed ways to find funds to purchase these computers.

They have been—unfortunately they have been limited in the amount of service that they can give university users because the computer itself is not the entire system. It requires a lot of peripheral devices associated with graphics, with networks, and so on. And these are quite expensive and have not been underwritten by the Federal Government to the present time and, in fact, limit the usefulness of these three centers.

We are, of course, expecting to have responses from these three centers to our solicitation for computer time, and we hope that we will be able to help these centers become generally suitable for the use of most scientific disciplines when they become centers for the allocation of time in our program.

Mr. Winn. What guidance, if any, is available for future prospective centers for these collective experiences? You said you weren't exactly sure.

Dr. Knapp. Well, I would certainly think that any future university center would discuss at some length the problems that the three centers have had with the people who, in fact, have put these together before they would take the plunge themselves.

As far as there being a cookbook or a guidesheet on how to do it for the next ones, I don't know of anything like this which is in preparation at the present time, but it could be done.

Mr. WINN. Do you think it might be done in your shop since you would be dealing with these people as being somewhat helpful to any other universities that might want to become a supercomputer center rather than just the age-old game of going to Uncle Sam for money?

Dr. KNAPP. As we go through the startup of our program, I am sure we will collect a good bit of this data and I would presume that it would be available in terms of minutes of meetings or perhaps could be pulled together as a document which could help future researchers.

Mr. WINN. Can you characterize their effectiveness to date?

Dr. KNAPP. I think all of these are still in the startup mode. I think there—maybe that's unfair. Maybe—I think they've been quite effective in certain areas of science but have not been generally effective over all of the disciplines and are just now becoming effective in the networking which they are doing privately, I believe, to outside university users.

Mr. WINN. Yesterday we were handed the Bardon report, which recommends spending over half a billion dollars in the next 3 years on supercomputers. With that kind of an expenditure, is everyone trained and ready, operators, researchers and everyone needed? And, if so, who would bear this expense?

Dr. KNAPP. I think that the specific numbers in the Bardon report are the result of a series of workshops and are not necessarily things that we could achieve at the present time.

As I said in my testimony, we are—Dr. Hayes, who is here with me today, is making plans for a NSF program to develop supercomputer access and has in it realistic—is trying to get in it with the input from the community realistic estimates of the number of trained people required, the number of trained people available, the rate at which these people can be trained and acquired in the universities for this and, in general, put together a realistic system.

I—the Bardon report is a report of a working group within the foundation and is being refined at the present time with a group which is planning our specific program.

Mr. WINN. What, in your opinion, would happen if the funding were reduced or deferred by 1 year or 2?

Dr. KNAPP. I think we would not initiate the program. See, we are proposing to spend—at 1984 money there was $15 million allocated for scientific—advanced scientific instrumentation. We are proposing to spend $9 million of that on advanced computing facilities. That $9 million is a start, but without substantially more funds in the future I think the program will really not even start to address the problem that Mr. Fuqua has mentioned of many computers for university use.

Mr. WINN. So you're thinking that most of these funds would be provided by the Federal Government or would you work out some kind of a cooperative effort with matching funds or something along that line with some of the States?

Dr. KNAPP. I think we would look at all the options that we could find: matching funds with States, matching funds with industries, industry participation, the possibility of multiagency funding of programs.

I think we are quite open to any possibilities and, in fact, our computing—our experimental computer research program, which has brought the medium-scale computing facilities at universities up in the last 3 years from almost nothing to now I think it is 16 or 18 universities have reasonable computing facilities in their computer science departments. That program has relied very heavily on shared funds with industry, and I would hope that we could do the same in the supercomputer area.

Mr. WINN. You presented the committee a program 1984—funding for 1984, 1985 and 1986 and it looks like your total on 1986 is $261 million.

Dr. KNAPP. That was the report of this—of this committee, the workshops that we had and is not the program of the Foundation.

Mr. WINN. It's not the program of the Foundation?

Dr. KNAPP. We are still developing the program.

Mr. WINN. You are still developing yours? I was going to ask you if OMB had been consulted and if they had agreed upon funding of that type.

Dr. KNAPP. No, we are still—that was—the report by Mr. Bardon is a report of a committee that met to discuss needs and did not—and did not have in it the realism of overall budgets of the Foundation and how we can work this program into our overall program with the budget constraints we know we will have.

Mr. WINN. Well, the Bardon committee didn't have to consult with OMB like you did, did they?

Dr. KNAPP. No.

Mr. WINN. I think that is pretty obvious.

Dr. Martin, to what extent is industry participating in the S-1 supercomputer development program at Lawrence Livermore Laboratory? Mainly to what extent is industry being advised of the program's progress?

Dr. MARTIN. Well, I am not sure that I can answer that. Certainly they're well aware of what is happening, and there is constant dialog between the people at Lawrence Livermore and university researchers and the active industry participants, so their work certainly is not going on in a vacuum, nor is the intermediate results of anything that they are learning and discovering going on in a vacuum.

But as far as formal interactions with industry, I would be glad to provide that information, but I cannot answer it right now.

Mr. WINN. I would hope that you could, because there seems to be a duplication or a redundancy and because the parallel processing design programs to be—to parallel similar industry effort.

I wonder if we are getting the maximum out of that program. And if you could furnish the committee any information, we will insert it in the record at this stage, as part of your answer to my question.

Thank you, Mr. Chairman.

Dr. MARTIN. I would be glad to do that. I think that I should, just as a matter of clarification, be sure that everyone understands that—that in a single processing unit we are reaching the technological limits.

Parallelism is a way of overcoming that obstacle, and that is putting a lot of different processors together side by side or configured in some way to make up for the fact that we can't exceed the speed of light. So, parallelism itself is just a way of overcoming a technological limitation.

And most people who are working in the computer area today in architectures would say that they are working with parallelisms, so it should not be misinterpreted that if people are looking at parallelism they are necessarily being redundant in the form of architecture that they are exploring. There are many alternatives. S-1 is one of them.

Mr. **Winn**. Thank you very much.

Thank you, Mr. Chairman.

[The information follows:]

From time to time industry is invited to status meetings and reviews of the program. I mentioned in my prepared testimony that more than 160 copies of the Structured Computer Aided Logic Design (SCALD) system has been released to industry. The next industry review meeting is being planned for March 1984. As more S-1 uni-processors become operational in building up the system, industry will be invited to participate in the evaluation of the system by running their programs on the actual equipment. The system will be put on the ARPANET, initially for conducting the government tests and evaluations. Plans are now being considered for making the system available to researchers in the academic community via the ARPANET.

The program is not redundant with industry efforts. No one else is trying to put together a number of very high performance uni-processors into a large multi-processor configuration. Parallelism in supercomputers will come in many forms. This is just one of many configurations with which we must experiment in order to understand how to best exploit the current state-of-the-art.

Mr. **Fuqua**. Mr. Skeen.

Mr. **Skeen**. Thank you, Mr. Chairman.

Dr. Knapp, I would like to ask you—you mentioned several times the cooperative atmosphere that's going on between the academic communities and industry. What is the—does NSF have any idea what the total national capability is now insofar as supercomputers, given what industry's participation is in its own sales or being a customer for supercomputers? What's the industrial side—what's the representation there? We've got three centers now I am given to understand in the supercomputer centers that are academically oriented. What's the industrial involvement?

Dr. **Knapp**. You mean how many computers are in industrial——

Mr. **Skeen**. Something that are—or what involvement? Say, for instance, the petroleum industry. What's their involvement in supercomputers? Has it been large, small or——

Dr. **Knapp**. It's recently grown, I think. In fact, the supercomputer that is installed at the Colorado State University campus near Denver in Greeley, I guess——

Mr. **Skeen**. Greeley? Fort Collins.

Dr. **Knapp**. Fort Collins? OK—is partly funded by the petroleum industry and will be used for oil field modeling as part of its use, its industrial use, even though it is located at a university.

The aerospace industry, of course, funded by the Department of Defense, has a rather major——

Mr. **Skeen**. It should be a natural.

Dr. **Knapp**. It's a natural for supercomputers, and there are quite a few installed at—in various—but as far as—I can find this out and I will submit it to you for the record.

Mr. **Skeen**. I was just very curious because it all boils down to how many people do we have that have access to the kind of training that they need. And I was curious as to what our total national capability could be as represented by industry or the academic communities and because I've had the feeling that we've been—we've reached some kind of a divergent. We are dependent on one another. Industry is dependent on how well you train these people in the academic centers and I note there is a close relationship and I just wondered how close it was—how closely it's being monitored, and how aware we are of what our supercomputer capabilities are as a nation.

[The information follows:]

Industry has become deeply involved in the use of supercomputers for oil exploration and recovery, electronic design, aerospace and automotive design, television and motion picture graphics, nuclear power, and computing service bureaus. Other applications are developing rapidly. A survey of installed systems in the U.S. at the end of 1983 is as follows:

Users:	*Number of machines*
Government research laboratory	23
Petroleum	6
Electronics	3
Aerospace	6
Automative	2
Nuclear power	2
Computing service bureaus	3
Motion pictures	1
University	3
Total	49

Dr. **Knapp**. I think that's one of the major reasons that we have to get supercomputers into universities and into being used by all of the scientific disciplines. It's a—it's interesting to think what kind of training is needed for someone to use a supercomputer say in industry. It's not only training in computer science. That's in fact probably—it's very important, but it's not the only kind of training that's necessary.

People who receive advanced scientific degrees in chemistry and physics and astronomy, in many different fields of the sciences, are the people who then go to industry and do the oil well—the oil field logging, the measurements, the simulations, the modeling of the atmosphere, say, at our NCAR facility in Boulder.

The training is very diverse for this field. It isn't that we need all these people just in the field of computer science. It's that we have to get these computers and their power through all of the sciences so that they can then be applied to all of the sciences at universities, but also in industry. And it's really quite a complex problem: how to get the training, how to get it stepped up so that more and more people understand how to use these machines.

Dr. **Martin**. Can I respond to that?

Mr. **Skeen**. Yes; I certainly would, Dr. Martin.

Dr. **Martin**. As a computer scientist, I frankly feel that is tackling the wrong part of the problem. You put that machine on the campus and you give them some money to maintain the machine and I will promise you you don't have a training problem.

The computer scientists and the engineers who can use that capability will dig in and roll up their sleeves and their intellectual motivations will surprise you to no end. And trying to brute force an educational program rather than just supply the tools that will allow these people to run is the wrong angle of the problem.

Mr. SKEEN. So, supply the—access to the computers isn't the answer to this thing? Is that what you are saying?

Dr. MARTIN. Access to the computers, and if you put the equipment on the campus, then you are going to have to supply the capability to maintain and house them. But if you put it there, it will be used, and the professors who are there will develop the courses necessary to disseminate the knowledge on how to use that machine. It is—it is not the bottleneck at the present time. The availability of the equipment is.

And if we could go one step further and find a way to create the economies of scale, you know, if you told the supercomputer manufacturers that they could place today's supercomputer—because tomorrow's hopefully will be tenfold or a hundredfold faster, but let's take this snapshot—they could take the supercomputer of today and put one on every single campus in the country and somehow we would make sure that the maintenance was taken care of, you could develop a program with industry, allowing industry to somehow write off underwriting some of that expense. And you would find a great explosion in the awareness of the academic community on how to use them.

Mr. SKEEN. Well, Dr. Knapp and Dr. Martin, that was exactly the thrust of the question that I had. And I was very curious—during the discussion—both of your discussions have been very good—just how dedicated industry is to the academic community or how well they're cooperating in making this kind of situation available to people that ultimately benefit them more than anyone.

Dr. MARTIN. They are looking for leadership, and there are avenues that can be opened. The R&D tax credits that have been offered to industry have helped tremendously, and I hope they will continue and enlarge over the years.

Allowing industry putting into place mechanisms by which they can provide to academia not only equipment but also support for that equipment would be another thing.

Mr. SKEEN. I think maintenance should be thrown in as an element, too.

Dr. MARTIN. Well, without it——

Mr. SKEEN. That's right.

Dr. MARTIN [continuing]. The university systems just cannot support the equipment. It's a fact of life.

Mr. SKEEN. I have one last question that I—that really concerns me, and it was brought up several times in your testimony. And that is the inability that we've found in the computer industry and in the sciences to keep up with the software section of the production.

I think it's alarming that we are at that point that we have the ability to produce the machines, but we're not doing the software work. And it's happening even in the kid's toy end of this whole business. They don't keep up with the software and now I see several industries—or several members or segments of the computer industry dropping out of some of the highly competitive end of it because they say they just simply can't keep up with the kinds of software that are required.

You mentioned one aspect of this that really had not—I had not been aware of and how labor intensive——

Dr. **Martin**. Intensive; yes.

Mr. **Skeen**. And can you give me some idea, just in a kind of a summary, of just how troublesome that part of the thing is?

Dr. **Martin**. Yes; and I can give you an analogy that was created by somebody else—and I don't know who it was, so I can't give them credit—but if we look back to the trend in the use of telephone systems in the early 1900's, there was a study done that said if then the trend that was observed over the preceding 5 years perpetuated itself for the next 20 years and they continued to operate the telephone system the way they were at that time, then it would take every woman in the United States to serve as a telephone operator.

Now, someone had the wisdom to realize that there were probably one or two women who did not want to be telephone operators and so they were probably going to have to find another solution. And they looked for a technological solution to the labor intensiveness of the approach they then had for providing this service to their customers, and the solution was the telephone switching system.

There were other things that happened, but the trend not only increased at the rate they had projected, it increased at a much faster rate. We have served a much larger user base and we've developed all sorts of extrapolations of the telephone. That may make some people happy and some not so happy.

Mr. **Skeen**. Then we have divestiture.

Dr. **Martin**. Yes, but we have an exact analogy in the computer area. What we have is a trend in increase in demand for software applications and we are seeing what ultimately could be an exponential growth curve taking us from—as you may recall, there were statistics that said we are spending $3 to $6 billion on software and hardware in 1980 and it would be $30 billion by 1990, which is 10 percent of the defense budget. Well, we can't spend 10 percent of the defense budget on the development and change necessary for software.

Now, the growth in that curve is comprised of a number of things. One is there would be more software because there are more applications for it. That's good. The other is there is a high cost of the software development because of the people resources. The reason we are disproportionately paying for the people there is because they are coming from a limited reservoir, and so you are paying a premium price for that talent.

The third component is that the method that you now have for developing that software is too labor intensive to meet this rapidly increasing demand. So, you have to find a way to bring things back in line and make that curve manageable and fundable, and that was the intent of our STARS program is, one, to expand the number of people who can be part of the computer science hardware and software resource, but also to hit at the base of the problem, and that is to make the generation and evolution of software less labor intensive so that we can continue to provide more applications to an ever-growing user base. It's important to us in defense, but it is important to the economy of the United States. We feel we can do that.

Now, why have we been able to do all these things in hardware and have not broken the back on the software problem? Well, at the time that we started out, all the cost was in hardware and very little was in software.

Mr. **Skeen**. That's where the interest was.

Dr. **Martin**. It was 1 percent or 2 percent of the cost of the system was in software, and that has changed dramatically and the situation has flipped.

We have not put into the software area the same intellectual energy, the same level of funding that we have in the hardware area. It's not that we're not smart enough. We simply haven't done it. And the problems that have arisen in the software area have been sort of insidious in their onset and now that they're here they're pervasive.

Well, we can't close our eyes and expect they will go away. They won't. We are going to have to apply those intellectual energies and we will have to fund the solutions. And if we do it, then we will derive the benefits. And if we don't, we lose not only the software but we lose the hardware industry to Japan and Europe and we also lose the keystone technology for our defense strategy of the future. So, we must do it.

Mr. **Skeen**. Thank you very much, Dr. Martin, Dr. Knapp. Dr. Martin, I might observe that I am certainly glad that some ladies today decide to become computer scientists, and we may need a whole bunch just like you, besides telephone operators.

Thank you very much, Mr. Chairman.

Mr. **Fuqua**. Thank you, Mr. Skeen.

Mr. Lujan.

Mr. **Lujan**. Thank you very much, Mr. Chairman.

I have always looked at our defense needs driving the computer technology. In your statement you say well: "We are through with hardware. Now the private sector can provide us with all that, and so we have to start concentrating on the software first."

It disturbs me a little bit that that's the kind of thinking because necessity has always been the mother of invention and because of our defense needs, we were able to move ahead in the supercomputers, but primarily because of our defense needs.

Is that going to change?

Dr. **Martin**. If that's the impression I gave, it wasn't totally accurate because, as I indicated, we fostered the development of the first large-scale computer in the forties. We funded the fabrication of the integrated circuits in the 1960's, and we are right now through my office supporting the VHSIC program, the intent of which was to bring the U.S. industry to a point of competitiveness in the world in the semiconductor area, and they were lagging behind.

What I intended to say is that we in defense look for the voids, that which is not being supplied by the commercial sector, and where we have demands that far exceed what they see in their commercial marketplace, and we push on those.

And what we have found in the supercomputer area is that we have—while we have not actually developed supercomputers, we have done the work or supported the work that was crucial to their existence in componentry.

Mr. LUJAN. Well, in a sense you did development because in the early years they were developed as they were needed.

Dr. MARTIN. OK. They were developed in response to requirements that we had, but many of them were developed by private companies, and that was what I meant. In the——

Mr. LUJAN. What I thought that you were saying was that, "OK, we have the capability, the hardware, for all the needs that we now have or foresee, so now we are going to turn it over to the universities and the private sector."

Dr. MARTIN. No, I think that supercomputing, as it is known today, should certainly be a commonplace thing in every university in the United States. We will still be pushing new frontiers, and we will have to, because what we need in our space defense systems is going to boggle the mind of the most visionary of us. So——

Mr. LUJAN. But the universities are good probably at supplying the manpower to develop them, to operate them, to do all of the different things, but they don't have the capability of creating the hardware, so to speak. So, it's two different roles, as I see them, and defense has got to play the continuing big role that it has.

Dr. MARTIN. As I did indicate in my testimony, NSA is working with supercomputer industry, looking toward the future, so we state many of the requirements. Those things that would normally be funded by those firms for expansion of their market should continue to be funded by those firms for the expansion of their market. Where we see that there is something we need that they would not otherwise be funding, that's where we jump in to fund them and motivate them, and they have been responsive to us.

Actually, I would say after the introduction of the integrated circuits industry ran with it. I mean, every microcomputer, hand-held calculator, multifunction watch and video game on the street is a fallout of defense research.

The commercial sector for 10 or 15 years just took off at such a rate that we really didn't have to push them a lot. Their energies were definitely providing new ideas at a sufficient rate to satisfy us.

It is right now that we are jumping back in and saying at the present time what you are doing to satisfy the commercial sector is not sufficient for what we envision as defense needs. That's why we have the VHSIC program. That's why over in DARPA we have the galium arsenide programs and their strategic computing program. So, when we have the results of those efforts.

Mr. LUJAN. Can I just interrupt you for a minute?

Dr. MARTIN. Yes.

When we have the results of those programs, we will apply them immediately to military applications. We have every reason to believe the commercial sector will pick up what we have contributed and use it in commercial applications, and so the supercomputers of 5 and 10 years from now are going to be hopefully an order of magnitude faster and better than what we have today, so we have the most challenging problems and we will stay out in front.

Mr. LUJAN. I have one more thing that I would like to visit with you a little bit.

One of the things that concerns me with our laboratories is that we just don't seem to make enough use of those capabilities that we have into the educational—bringing them into the educational system.

As a matter of fact, I got into a little trouble over at Sandia the other day, both with Sandia and the university, because it seemed to me like having both Los Alamos and Sandia so close to the University of New Mexico, that if somebody was going to study say nuclear engineering, for example, that the first university they ought to look at if there was that good cooperation would have been—would have been the university, because of the existence of the two laboratories.

The same perhaps applies in the subject that we are discussing now. We are talking about putting all of these computers into the universities, and perhaps we are overlooking that great resource that we have there. We have the machines—and I don't know how somebody could hook up to them or what use of the machines themselves the universities could make—but also the personnel as instructors. I didn't see that in your testimony.

Is there some way that we can take advantage of that resource that we have?

Dr. KNAPP. Well, I think that universities are taking advantage of that resource, at least at Los Alamos. I'm not sure about Sandia or Livermore.

Mr. LUJAN. They are. I am just saying, you know, that it just doesn't seem to be enough.

Dr. KNAPP. The Los Alamos laboratory has a network with I think 1,000 users around the country which come in and use their supercomputer facilities on programs which are of joint interest to the DOE and the university researcher.

A requirement on that program is that they be doing joint research with someone who is funded by the Department of Energy. I think that's a reasonable requirement considering the purpose for which those machines are at Los Alamos, and I think there is a fairly active program with the University of New Mexico, at least in the areas that I am aware of, in your State.

However, I think that in the long run the program that we have to develop in the Science Foundation for academic access to supercomputers involves computers at the universities and not in the major national laboratories.

I think it's important to have the computers where students can work on the operating systems, they can work on the sort of insides of these things at least part of the time, and that there not be—that it really be part of—the computer itself be part of the academic environment.

Mr. LUJAN. There's not enough computer time to do more—I'm thinking even on a basic level, not necessarily research programs that the universities may have, but even just from an instructional point of view. If you want to be a mechanic, you got to get your hands on that motor at some point, and I assume that if you want to be a computer scientist you have to get your hands on that computer.

I thought that having these super—not in terms of supercomputers but super facilities, that a student just being around them will absorb that—will have hands-on experience with it. I just thought NSF might be able to kind of work together with those laboratories.

Dr. **Knapp**. We are working with DOE, but it's—they have mission requirements, too, and that's one of the problems that we see.

Dr. **Martin**. That was precisely it. If, in fact, DOE or DOD diverted too far from their specific mission, I'm sure that the Senate and the House would be the first to bring it to our attention.

Mr. **Lujan**. No; but you're saying that the greatest mission is to have those people being able to operate those computers, so there's a mission, also.

Dr. **Martin**. If you would like to augment our mission statement, we'll welcome that. Could we please have an increase in TOA as well?

Mr. **Lujan**. Thank you, Mr. Chairman.

Mr. **Fuqua**. Thank you, Mr. Lujan.

The committee will stand in recess. We have a vote. We will resume as soon as this vote can be completed.

[Recess.]

Mr. **MacKay** [presiding]. I would like to resume the meeting.

We may have other members come back. There were members here when the vote was sounded, and they have not yet had a chance to ask questions. Pending their return, I have a question of Dr. Knapp.

Dr. Knapp, yesterday two witnesses commented about the NSF and supercomputers, and both of them commented that they had concerns that there might be a structural problem that prevented NSF from really dealing aggressively with the supercomputer issue.

Now, they specifically were talking in terms of establishing centers. Dr. Kozmetsky talked about perhaps as many as 20 computing centers around the country, and he set forth a need not only for better state of the art in academic research, but also for having the equipment available to train people. He felt that if we were not able to do that, we ran a great risk of a delay which could cause us to lose our potential for market share.

I interpreted what he was saying to mean that NSF is structured along discipline lines and that these are, by their nature, mostly disciplinary instruments and that there might not be a mechanism within NSF to deal with that.

I would like your comments on that.

Dr. **Knapp**. I think that the program that we outlined and that we are developing does address some of those problems and, in fact, is a good solution to them. We expect to have the time allocated by the disciplines for various scientific research.

We expect to have the centers run by the universities reporting to another organization within the foundation which would be responsible in some way or another, which has yet to be determined and partly depends upon the scale of the response we have and the kind of program we are planning for, the centers to be operated with funding from a centralized source within the foundation. I feel that that—it's essential to have both of those kinds of responsibilities within the foundation and located where they would be located.

I realize that the use of the computing facilities is multi and interdisciplinary, and having located at one site at a university research in various disciplines is a very attractive option which we really wouldn't have much control over from the foundation. That's the university's choice.

There has been a very interesting proposal from Cornell University by Mr. Wilson and others who have testified to you about a possible theory institute at Cornell which would have under its own roof, with a computer as a major tool, disciplines as diverse as psychology and mathematics and physics and chemistry, so that these disciplines could share experiences, tools, software difficulties with one another. I think that's a good model myself, but that's really the university's choice, not the foundation's. But we will within the foundation certainly be a centralized location for funding the centers of computers themselves.

Mr. MacKay. Dr. Knapp, what is your timeframe for implementing this plan?

Dr. Knapp. We are going to purchase the computer time for disciplinary use on existing supercomputers or supercomputer at supercomputer centers this year, 1984. We are trying to move as quickly as we possibly can, and in 1985, depending on the program as it evolves, we would probably move toward new university-based centers.

Mr. MacKay. All right. I have no further questions.

Mr. Carney.

Mr. Carney. I have no questions, Mr. Chairman.

Mr. MacKay. Thank you very much, Dr. Knapp and Dr. Martin, for your participation.

We will—without objection we will file the testimony of Dr. Keyworth, and since he will not be available for questioning, we will have his testimony, his written testimony, as part of this hearing.

Mr. MacKay. If there is no further question or business, this meeting is concluded. Thank you.

[Whereupon, at 12:05 p.m. the committee was adjourned.]

APPENDIX

Additional Testimony Submitted for the Record

STATEMENT OF
DR. JOHN BRADEMAS
PRESIDENT, NEW YORK UNIVERSITY

BEFORE THE
HOUSE SCIENCE AND TECHNOLOGY COMMITTEE
NOVEMBER 15, 1983

I submit this statement today on behalf of the New York University Courant Institute of Mathematical Sciences, which is considered one of the outstanding mathematics institutes in the United States. My colleagues and I congratulate the Members of this distinguished Committee on holding these important hearings on supercomputer research and on the appropriate roles that each sector of society must play in the financing and developing of supercomputers.

Our statement is intended to supplement that being presented to you by the Association of American Universities and the National Association of State Universities and Land Grant Colleges, which we support. The purpose of our testimony is to highlight for the Committee two specific issues.

First, there has been considerable attention given to the question of which Federal departments should have a role in supercomputer research and development. We believe that it is imperative that several Federal departments continue to play a vigorous role in developing supercomputers and that funding for this purpose be increased. We wish in particular to discuss the historical position and potential future participation in such activities of the Department of Energy, the Department of Defense and the National Science Foundation.

Historically, DOE and DOD have supported research and development for supercomputers. The major-- but not only-- interest of DOE, for example, has been the weapons laboratories, principally Los Alamos and Livermore. In the past five years, however, these laboratories have relied their major computations on the very successful CRAY machines.

In the last year or so, several factors have revived, at least for the time being, an interest in new types of machine architecture:

(1) The "Very Large Scale Integration" revolution has suggested the possibility of many imaginative new approaches to the design of supercomputers.

(2) Computers are now perceived to undertake many more activities (such as interpretation of military communications in "natural languages") than in the past.

(3) Japan is apparently determined to become a world leader in computers.

Acting through DARPA, DOD has decided to resume supporting the development of new computer architecture. DARPA has chosen a bold and innovative role, the need to develop powerful and advanced computers for artificial intelligence (equivalent to the Japanese Fifth Generation project), arguing that the task of supporting new machines for "scientific", i.e. numerical, processing was being done elsewhere. This assertion was not, except for the work at CRAY and Control Data Corporation accurate. CRAY and CDC have naturally focused entirely on logical continution of their vector machine designs.

The Administration has recently moved to fill this gap and has directed the DOE to act as lead agency for numerical processors. It is logical that

DOE should also be involved because its laboratories require computational power for reactor and weapons design, materials science, accelerator design, and the analysis of high energy physics experiments. But Congress has not yet authorized DOE to enter this field and even if Congress decides to do so, DOE's entry is not likely to begin before FY '85. We believe that such a DOE program is important and should begin in the coming year. DOE has historically been a steady supporter of research in computing when there were few other sources of support, and the Department has given long and stable support to our work at Courant.

The National Science Foundation must also be involved. NSF is now considering large scale support for access to contemporary supercomputers by university researchers. Such access will have the important consequences of greatly improving the quality of basic research, of training students for the style of research of the future in which computing will play a very large role, and of strengthening the market for large machines. This focus on "access" does not, however, directly address the need to investigate new ideas in the organization of computers. NSF, in our view, must play a strong and expanded role in funding basic research in this field. We have suggested, for example, that one appropriate role for NSF would be to support advanced computer science centers where different models of the supercomputer could be "test-bedded".

New York University has what we believe to be a highly promising design, the Ultracomputer. The Ultracomputer appears able to bridge the gap between numerical and non-numerical computers, and has therefore attracted interest from both DOD and DOE. We believe, further, that these two kinds of computing are not as antithetical as is often supposed-- "expert systems" often have to deal with numerical data, or even to direct the course of very large numerical calculations. Equally important, a general purpose machine of this kind will give researchers the freedom they must have to explore new ideas but will also be most attractive to future commercial exploitation.

Because this area is so important for the technological and economic health, and the defense of our country, we hope that you will be able to encourage the support of both Congress and the Administration for a program of research and supercomputers in general, with appropriate balance among DOD, DOE, and the NSF.

Second, we at New York University are concerned that the role envisioned for universities in this issue may be too narrow. Whenever university participation is discussed, it is often in connection with the difficult question of access to supercomputer capacity. The National Science Foundation and the university are all too often, in our view, left as their assignment, solving the access and networking questions. We would strongly suggest that the university and the Federal government should and can play a role in the development of tomorrow's supercomputer capacity. One of the final conclusions of the Lax Panel was that any national program include a focus on producing adequate research and development capacity basic to the design and implementation of new supercomputer systems of substantially increased capability and capacity beyond those likely to arise from commercial sources. We agree that that research and development cannot be left to the private sector alone. For this reason we not only urge a greater Federal investment, but also suggest that the university can play a critical role in the development area as well. University-based research is not subject to the same set of constraints which confront the private sector -- the concern over marketability and economic return to the company.

Allow me to make the following five points in support of this view:

(1) The future development of computer technology in the United States requires cooperation among universities, industry and the government. Industry has a strong vested interest in "vector processing", and needs incentives to broaden the prospects of new architectures. Venture capital interests have a two year horizon, and because research and development time for supercomputers is likely to be longer term, venture capital is not likely to be a source of funds.

(2) Many imaginative designs for supercomputers are starting to emerge from university laboratories. We must find a way to support such efforts.

(3) Universities also play an important role as a source of future researchers, computer architects, language and software designers as well as users and developers of new appliction areas and methods.

(4) There are several exciting projects now underway, some activity involving both industry and university participation. The success of these projects will eventually be tested in the marketplace, and will probably involve both existing as well as future companies.

(5) We believe that "scientific" (i.e. numerical), machines, and "thinking" machines are not antithetical. A general purpose computer well suited to both applications is possible. If such a machine is built, it will be attractive not only researchers for the freedom it will give them to explore new ideas but will also be attractive to commercial developers as it will give them the widest possible market.

We appreciate this opportunity to present our views. I look forward to working with you on this issue as you address the need for a national program in FY '84 and FY '85.

James C. Browne

October 13, 1983

James C. Browne was born in Conway, Arkansas, on January 16, 1935. He received the B. A. degree from Hendrix College, Conway, Arkansas, in 1956 and the Ph. D. degree in physical chemistry from the University of Texas at Austin in 1960.

He worked from 1960 to 1964 as an Assistant Professor of Physics, University of Texas at Austin, carrying out quantum mechanical calculations of small molecules, and doing research and development work on operating systems. He was, from 1964 to 1965, at the Queen's University of Belfast, Belfast, Northern Ireland, on a National Science Foundation Postdoctoral Fellowship. He was, from 1965 to 1968, Professor of Computer Science and Director of the Computation Laboratory at the Queen's University of Belfast. In 1968 he rejoined the University of Texas at Austin as Professor of Computer Sciences and Physics, and served as Chairman of the Department of Computer Sciences in 1968/69 and from September 1971 until January 1975. His research interests include operating systems, systems modeling and design, performance evaluation of computer systems, and theoretical computations of molecular properties. He has published about 125 refereed papers on technical subjects.

Dr. Browne is a Fellow of the American Physical Society and the British Computer Society, and a member of the Association of Computing Machinery.

COST EFFECTIVE APPLICATION OF SUPERCOMPUTERS

J. C. Browne
The University of Texas at Austin
Austin, Texas 78712

October 24, 1983

Table of Contents

1. EXECUTIVE SUMMARY
2. APPLICATION DEVELOPMENT AND CURRENT AND FUTURE GENERATIONS OF SUPERCOMPUTERS
3. A MORE VERSATILE GENERATION OF SUPERCOMPUTERS
4. COST REDUCTION OF SUPERCOMPUTER SERVICES
5. SYSTEM ARCHITECTURE FOR THE NEXT GENERATION OF SUPERCOMPUTERS
6. INDUSTRY/UNIVERSITY/GOVERNMENT COLLABORATION ON THE PROTOTYPING AND EVALUATION OF RESEARCH CONCEPTS OF FUTURE SUPERCOMPUTERS

EXECUTIVE SUMMARY

The economic benefits of supercomputers come not from their mere existence but from their *application* to those important problems which require (or can cost-effectively utilize) very large units of computer service. Current applications of supercomputers are largely limited to numerical modeling of the behavior of physical systems such as the earth's atmosphere (weather forecasting), movement of viscous fluids through porous media (oil field reservoir modeling) and similar applications. The potential of such numerical simulations for their impact on both science and engineering and the economy has hardly been tapped. It is the case, however, that extension of computer-based modeling to other complex systems such as production systems (factories) or economic systems offers potential return which far outscales the current class of applications. The potential impact of intelligent systems which are operational in real-time (intelligent robots such as missile guidance systems) is even greater yet. These applications potentially impact productive industry and commerce to a greater extent than service industries are being impacted by use of small units of computer power through microcomputers. These potential applications of supercomputers to new types of applications require system architectures which can deliver very high performance computing on numeric processing, symbolic processing and signal processing.

There are both cost and technology barriers currently precluding attainment of these benefits. Application of the current generation of supercomputers to these new problem domains is awkward and thus often not cost effective. The awkwardness arises from the absence of software support for these applications and because the architectures of current day supercomputers is specialized to most effectively service numerical computation. Awkwardness of use leads directly to high cost of application. The span of applications which can be reached with today's supercomputers, while very far from fully explored, is too small to give the cost reductions obtainable from large-scale production. It is further the case that the computer power attainable from systems built according to the design principles of today's supercomputers will not extend to meet these application requirements. A new generation of technology is required. The next generation of technology is currently very immature. We cannot yet effectively select among the many opportunities it offers nor, of course, do we know how to apply the resulting system architectures to the new applications.

The role which the government can play in developing the potential of supercomputers is to support research and development which will breach these cost and technology barriers. The problems to be solved are:

- development of the applications of very high performance computing
- development of future generations of supercomputers with the versatility to attack new applications
- development of a technology basis for orders of magnitude cost reduction in the delivery of large units of computer service

The recommendations for government action to *assist* in solving these research and development problems are:

- Establish easy and cost-effective access to supercomputers for university researchers and high technology small business to explore *new* applications of supercomputers. This access must include not only current day supercomputers but also examples of the innovative architectures likely to be the basis for future generations of commercial supercomputers.
- Establish a research and development program for development of a generation of *versatile* supercomputers capable of cost-effective application to the economically significant problems.

- Establish a research and development program which focuses on the development of software and applications for new generations of supercomputers. This program can be combined with the program for access to supercomputers given as the first recommendation.
- Establish mechanisms for assisting in the movement of research concepts for future supercomputers to the marketplace.

APPLICATION DEVELOPMENT AND CURRENT AND FUTURE GENERATIONS OF SUPERCOMPUTERS

The economic importance of supercomputers will be determined by the range of applications which they *cost-effectively* support. Thus the governing factor for the utility of a given supercomputer is the range of applications it can cost-effectively support. It must be recognized that cost-effectiveness is impossible to assess in the realm of applications which develop new knowledge or in the realm of national security applications such as weapons design.

There are two theses for this section:

- Applicability of even the current generation of supercomputers is not being exploited and should be.
- There can be developed future generations of supercomputers whose cost-effective range of applications can potentially revolutionize American industrial productivity. The development of these supercomputers will be motivated by perception of the markets opened up by development of prototypical new applications.

The attainment of goals implicit in these theses and recommendations for government actions to stimulate the market forces to attain these goals will be given.

The selection of applications for development is based on two factors: technological feasibility and return on investment of time and capital. Applications of supercomputers are not proceeding at their potential because

- Access to supercomputers by university research groups and high technology small business is awkward.
- The computing environments of current-day supercomputers do not commercially support a broad range of applications.
- The development of new applications of supercomputers generally requires large capital and effort investments and the returns on investment may be long in coming.

These are the problems which must be overcome in order to stimulate the development of a generation of applications of supercomputers which will realize their vast impact on the productive economy. The fundamental role of the government is to create an environment where these opportunities for applications of supercomputers (with their great national significance) will become attractive to the free market entrepreneurs. The current situation is that the *long term* national interest of the United States and the *short term* interest of the entrepreneurs do not coincide.

A key factor is that there are far more opportunities in the burgeoning world of applications of computers than there are qualified exploiters of the opportunities. Thus there will be a selection of small and easy projects which maximize short term return. It is extremely easy and inexpensive for a computer scientist and a venture capitalist to develop and market a package providing a personal service function on a microcomputer. Complete access to a microcomputer is a matter of a few thousand dollars, the development and marketing cost is measured in thousands of dollars, the return, if the product is successful, is immediate, there is now established a marketplace mechanism for these products, the unit cost of the product is small and there is a large market. *None of these conditions holds* for applications of supercomputers to large-scale problems. The applications of supercomputers will lag until some or all of these circumstances are altered. The first and most important barrier is *access.* This is something the government can easily facilitate.

RECOMMENDATION

Access to supercomputers should be established for university researchers and high technology small business concerns for the development of new applications of supercomputers. This access should include not only state-of-the-art current day supercomputers but also the innovative architectures which will be the basis for future generations of supercomputers. Only if this latter portion of the recommendation is followed can we hope to begin establishing a large base of applications to drive development of future generations of supercomputers.

A MORE VERSATILE GENERATION OF SUPERCOMPUTERS

The current generation of supercomputers is primarily specialized to be efficient on large-scale numerical simulations of physical phenomena. This specialization is reflected in both hardware and software. The result is a fairly severe limitation on the range of problems to which these systems can be easily applied. This limitation is a result of the fact that the *past* market for high performance computing was entirely in numerical simulation applications. The diversity of the potential new applications which stretch out before us now calls for much more versatile system architectures which can be cost-effectively applied to these new applications. The hallmark of these new applications is that they combine requirements for high bandwidth interfacing with the real world, numerical simulation of complex physical systems and automated reasoning. (Interfacing of computers to signals from their environment is called signal processing. The processing associated with automated reasoning is commonly called symbolic processing. Each type or mode of processing has somewhat different characteristics.) Future generations of supercomputers must cost-effectively meet all of these requirements if these opportunities are to be accepted. The next generation of supercomputers being planned by the industry and being planned by mission-oriented government research and development programs appears to be perpetuating this unfortunate specialization. Industry is focusing on more powerful numeric processors while the military is focusing on the development of supercomputers for symbolic processing to implement decision support systems. We very strongly believe that the *economically important* applications of supercomputers will require super computation power in all modes. There is further synergism in the power to be attained from each mode. Effective use of automated reasoning based on symbolic processing can guide numeric processing to follow near optimal solution paths. Numeric processing can be substituted for symbolic processing where special case solutions are needed because of the explosive growth of possible general solutions based on symbolic processing. Recognition of the effectiveness of this synergism is just beginning to emerge. The first studies of the potential of this synergism are just now beginning to be initiated. It is significant to note that this synergism is being explored by *application developers not* by the protagonists of each class of computing.

RECOMMENDATION

The government should establish a research and development program focusing on the development of versatile supercomputer systems which can cost-effectively and synergistically implement very high performance computing in all modes. This research and development program should proceed from study of applications of such versatile supercomputers working toward the establishment of a market for such supercomputers which will stimulate private industry to extend its products to incorporate such capabilities in future generations of supercomputers.

COST REDUCTION OF SUPERCOMPUTER SERVICES

Cost-effectiveness is the ultimate determining factor for the applicability of supercomputers. It is essential that the total cost for delivery of supercomputer services be lowered by an order of magnitude if very high performance computing is to be broadly applied. The total cost of an application system includes:

- development cost for the supercomputer system itself
- development cost for the application
- replication costs for the supercomputer system
- operating costs for the system

There are many avenues to cost reduction. A major one is increasing the breadth of applications thus increasing production volumes and taking advantage of economies of scale. The avenues to increasing applicability were addressed in the preceding sections on applications and more versatile system architectures.

Another avenue is to reduce the cost of the components from which the systems are constructed. The current thrust of component technology is to support very low cost for low and moderate performance system architectures. These dominate the current component marketplace. (The VHSIC program of the Department of Defense is an exception.) There is little incentive for component vendors to develop very high performance components since the current marketplace is small. Stimulation of research and development for very high component technology is needed in order to lower the cost. This could be a systems oriented extension of the VHSIC program.

RECOMMENDATION

There should be established a research program aimed to providing a technology base leading to order of magnitude cost reductions for development and replication of very high performance computers.

SYSTEM ARCHITECTURE FOR THE NEXT GENERATION OF SUPERCOMPUTERS

It is an accepted precept in the industry that the next major generation of supercomputers will involve multiprocessor configurations and parallel structuring. The impact of explicit parallelism on software and applications will be enormous. There is an almost total lack of detailed knowledge of parallel problem structuring and the software which will be necessary for parallel computation. The availability of algorithms and software which support efficient development and effective execution of programs on future parallel supercomputers will therefore set the pace for the application of the next generation of supercomputers. It is essential that for application of supercomputers there be available basic methods and software to support formulation and development of applications on future generations of parallel architectures.

Thus development of future generations of supercomputers based on parallel computing is a systems problem not merely a problem in developing computer hardware. Failure to recognize this salient fact will substantially postpone or diminish the leverage of development of commercially usable high performance computers based on parallel computing concepts.

Development of basic software for future generations of architectures will be intrinsically more difficult than the development of basic software for existing supercomputer systems. This is because the present generation of supercomputer systems has been developed by evolution from current mainframe architectures. Thus the concepts for applications and software have been obtained by evaluation from past mainframe systems. This will not be the case for the software systems and applications of a future generation of supercomputers based on parallel structuring. If we are to exploit future generations of supercomputers in a timely manner we should begin now the development of the basis for their application. If we have not learned from past experience the folly of deferring application software issues to follow architectures, then we will suffer the penalty of repeating this avoidable mistake.

RECOMMENDATION

The first RECOMMENDATION to provide access to supercomputers for development of new applications and scientific research should be extended to explicitly include innovative parallel architectures for development of software and methods for applications of parallel computing. Explicit funding of research and development of software and applications for parallel architectures should be emphasized.

INDUSTRY/UNIVERSITY/GOVERNMENT COLLABORATION ON THE PROTOTYPING AND EVALUATION OF RESEARCH CONCEPTS OF FUTURE SUPERCOMPUTERS

Many important concepts for the future generations of supercomputers are now being developed in university and other research laboratories. The cost of prototyping research concepts for supercomputer systems greatly exceeds most current research program budgets. Thus the practical evaluation of these research concepts and the transfer of technology from research to development to products is delayed because the evaluation cannot be effectively accomplished. There is no mechanism for sorting out from the current plethora of concepts those which are most likely to develop into effective products. Industry does not have procedures for selection of alternatives and individual companies may not be able to afford to take a comprehensive approach since the costs of prototyping and evaluation are very high. It is clear that industry participation in the prototyping and evaluation process is critical. The target of the evaluation is to sort out "the chaff from the wheat" so that private enterprise and venture capital will pick up the attractive products for development and commercial exploitation.

One useful structure which is arising from the Japanese model for developing "supercomputers" and "Fifth Generation" architectures is the establishment of a critical mass for the prototyping and evaluation of important concepts in super computation in two vital application areas.

Therefore the current situation leads to circumstances where research concepts currently being exploited are likely to be evaluated and ultimately exploited earlier overseas than in the U.S. government.

RECOMMENDATION

A program of university/industry/government collaboration for prototyping and evaluation of supercomputer concepts should be established. This program should be founded on the traditional mechanisms of research and development sponsorship but should incorporate university/industry collaboration.

Statement

of

Dr. John Gibbons
Director
Office of Technology Assessment

INTRODUCTION

Thank you for asking me to testify today on the subject of supercomputers. I would like to have my formal testimony placed in the record and present here a summary of it in my oral comments.

OTA has been looking at the state of U.S. computer technology in a number of studies; these include our soon to be published report on competitiveness in the microelectronics industry and our ongoing assessment of the state of research and development in information technology. That latter study was requested by this Committee, as well as the House Committee on Energy and Commerce. In the course of our R&D study, we convened a workshop in July on the subject of advanced computer architecture.

My comments today will be based on this work. I must point out, however, that the R&D study is not complete, so my observations should be taken as preliminary.

In short, supercomputers are the largest, most powerful computers currently available.

But that definition begs many questions, the most important being "the most powerful for what?" The high performance of a supercomputer is obtained, in part, by optimizing its architecture (or design) to solve particular classes of problems. While past supercomputers have been designed for large-scale mathematical computations, we now see computer science research laboratories are now working on a greater variety of designs intended to address different applications. For example, the Japanese "Fifth Generation Project," the British "Alvey" program, and our own DARPA program on advanced computer architecture are all aimed at artificial intelligence applications.

This optimization of design to applications explains why the performance of a supercomputer is highly sensitive to the type of work we give it to do, an important point that relates to plans to improve access by researchers.

Another important point to keep in mind is that the supercomputer is a special case of a broader class of computers that use "advanced architectures." Not all of these machines have price tags in the several millions of dollars, yet some of them offer performance for very specific types of problems that nearly approaches that of a supercomputer. While the subject of the hearings is supercomputers, much of what I have to say in my testimony also relates to this broader class of machines.

FINDINGS

By and large, OTA finds the U.S. computer industry to be very healthy. It is the world-wide leader in technology and in sales, and the prospects are good for continued leadership in the forseeable future are good. Until very recently, U.S. firms have been the only vendors of supercomputers. Now, new competitors have appeared, and they are determined to at least capture part of a market that appears to be expanding rapidly away from its original narrow base of scientific and defense users. To the extent that competition accelerates the pace of development and brings more powerful technology to market faster, it may be welcomed by the scientific users who are constantly pushing the limits of available computer power.

At the same time, we should not be complacent. Advanced computer architecture is a fundamentally important field of research. It results not only in new supercomputers, but in design concepts that eventually work all the way down the line from commercial mainframe computers to mini and micro computer systems. Computer architecture researchers and the supercomputer user community both face impediments and problems that this committee can play a role in easing.

POLICY ISSUES

This Committee faces two basic policy issues: providing access for basic researchers to supercomputer capability, and assuring an appropriate level and direction of research to support future advancements in supercomputer architecture.

The National Science Foundation has done an excellent job, through both the Lax Committee report and the report of their recent working group, in laying out for us the importance of supercomputers to basic research in all fields. Applications abound: For example, in some fields of science, supercomputers solve large systems of mathematical equations that represent physical processes such as the formation of a star or galaxy, the behavior of the plasma in a fusion reactor, or the chemical properties of a complex molecule.

There seems to be little doubt that access to supercomputers by U.S. scientists is woefully inadequate. Only three research universities in the U.S. currently operate supercomputers. Providing for the needs of researchers requires a three-pronged approach: new hardware resources, improved remote network access to the systems, and software development to support specialized new applications in particular disciplines. The NSF plan outlined by Dr. Bardon wisely provides for all three.

I would like to emphasize a few points. In order to develop real centers of expertise, the facilities need to focus on areas of research that share similar problems and research approaches. Furthermore, they need to be administered in such a way that they are used only for research that could otherwise not be undertaken. Such research would be characterized by a need for hundreds of hours, not seconds or minutes of computer time per year. It would be carefully tailored to the particular strengths of the machine and available software. Management of the facility and allocation of its resources might resemble that of an accelerator or telescope more than that of a traditional research computer center. Finally, assurance of long-term support and investment in adequate software and consulting staff is vital to encouraging the development of significant research programs.

The second major policy question is the state of support for research in new computer design.

The issue is particularly important with supercomputers, for the U.S. manufacturers of these machines are, by and large, very small firms. They commit their full resources to designing the next generation machine and, hence, have little funds for long-term, basic research. The Federal government has supported development of next generation systems indirectly by serving as the major customer, in a sense, guaranteeing a market for them. The case will remain the same for the next few years, although as a commercial market grows, Federal influence will likely diminish.

The situation is quite different with respect to the basic computer science research that will underlie succeeding generations of computers. Much of that has come and will continue to come from university laboratories supported by DoD, NSF, other Federal agencies, and private industry.

Multiple sources of support for research has long been considered to be a healthy characteristic of U.S. science policy, and it would seem to be so in this case. A variety of different, competing architectural concepts are being investigated now in U.S. research laboratories, many of which seem worthy of exploring further.

The main problem with many basic research projects in computer architecture is that they stop far short of implementation. Paper and pencil or "toy" implementations are not sufficient to adequately explore or demonstrate the potential of a new concept. In particular, there is little incentive for software and systems researchers to work on problems for a theoretical machine. Yet research is critical to understanding the true potential of a new design. Furthermore, because the research stops short of systems design, computer science and engineering students were not trained to think in those terms, decreasing their usefulness to industry.

CONCLUSION

The issue is important. A commercial marketplace for supercomputers is growing rapidly at the same time that our industry is facing foreign competition for the first time. Nonetheless, it is important to recall that not all research benefits from use of supercomputers, and not all computer science research benefits from supercomputer research. It could be a serious mistake to divert resources such as funds and research manpower away from other promising fields of investigation simply because supercomputers are a glamorous topic of the day or as an automatic response to Japanese or European plans. To keep priorities in order, the entire scientific community needs to be involved closely in the decision-making on allocation.

I am pleased to have been a part of this Committee's examination of the state of our science policy as it relates to support of advanced computer architecture, and would be pleased to answer any questions.

STATEMENT FOR THE RECORD

of
Donna L. Valtri
Project Director

OFFICE OF TECHNOLOGY ASSESSMENT
UNITED STATES CONGRESS

BEFORE THE

COMMITTEE ON SCIENCE AND TECHNOLOGY
U.S. HOUSE OF REPRESENTATIVES

SUPERCOMPUTERS

November 28, 1983

This Statement is based upon the findings of a workshop OTA held this summer on advanced computer architectures as well as additional study in this area undertaken during the course of OTA's project on Information Technology Research and Development. The research in this project has not yet been completed; therefore, this Statement represents only OTA's best understanding of the issues involved in this area to date.

Three of the major issues raised in the OTA workshop are addressed below. First, some of the major areas of scientific endeavor and industrial activity that are now or in the near future will be dependent upon the availability of supercomputers are discussed. Also highlighted are several areas of basic research that are likely to require, or would be enhanced by, the availability of supercomputing power. Second, the concerns of university scientists, manufacturers of supercomputers, and Federal agencies (both as purchasers of supercomputers and as funders of supercomputer research) about access to current and future generations of supercomputing machines are addressed. The third area of discussion is the state of U.S. research and development in supercomputers.

Fields of Research That Are Limited by a Lack of Computer Capacity and/or Speed of Computation

Exploration of problems in a wide variety of scientific disciplines and in engineering is dependent upon the availability of fast computers with large memories. The current generation of supercomputers takes advantage of the regular, repetitive nature of many scientific and engineering calculations. The machines overlap the processing of elements of arrays and other highly structured sets of data. This overlapping of computation affords the supercomputer its speed and distinguishes it from other computers, including most commercially available mainframes.

The manner in which the hardware of a supercomputer processes data is dependent upon what is referred to as the machine's architecture. Supercomputer architectures are highly specialized and finely tuned to the exacting specifications of certain scientific and engineering problems. The architecture built into today's supercomputers enables them to solve a variety of problems including (1) modeling and simulation of complex systems and events, (2) computer-aided design, and (3) complex real-time operations.

The first category, modeling and simulation, represents the most typical current use of supercomputers. There are several examples of problems which can only be studied through simulation because actual experiments are either too costly, or too dangerous. Nuclear weapons research has both of these characteristics and is currently the major modeling and simulation application of supercomputers. In addition, problems of fusion energy, aerodynamics, oil exploration and reservoir management, and communications networks are now, or are expected to become modeling and simulation applications of supercomputers.

In oil exploration, the models of geological formations possible without supercomputers provide only rough estimates of the recoverable deposits a field may contain. More accurate modeling with supercomputers will improve the likelihood of finding and recovering commercially salable quantities of deposits and reduce exploration investments. Several major oil corporations, including Chevron, Arco, Exxon, Texaco and Sohio have purchased Cray 1 or Cyber 205 machines. In addition to geological modeling, oil companies use their supercomputers to optimize the types and quantities of refined products from crude supplies.

Currently available computer simulations of weather phenomena do not account for major features of the landscape, such as mountain ranges, that can be critical to weather prediction. The ability to accurately model weather, which may be made possible by supercomputers, could enable researchers to make more precise forecasts for agricultural and other uses, and provide improved information about impending weather disasters.

One major design application of present fast computers is very large scale integrated (VLSI) circuits. The design of these circuits requires enormous computing power for the generation of high resolution graphics. Bell Telephone Laboratories devotes most of its Cray 1 to this use. Another major application in this category is airframe design. Boeing and NASA-Ames use their supercomputers for this purpose. In addition, the entertainment industry is making use of supercomputers in the design of animated films and cartoons and in the production of special effects. In the future, design of air and space craft propulsion systems, automobiles, and other complex tasks that can benefit from computer-aided design will require the use of supercomputers.

Researchers at General Motors have a mainframe-based computer-aided engineering system which minimizes the mass of a part while maintaining its strength. Similarly, Ford can model the performance of a car body or engine part without actually building it; an engineer can design and redesign a part many times before a final version is decided upon. As computer-aided design systems like these become more accurate and are applied to more intricate parts and processes, their computing requirements will reach supercomputer levels.

The third applications category, complex real-time operations, may eventually constitute the largest market for supercomputer power. Avionics (aircraft control and instrumentation), robot vision and control, speech recognition, adaptive control for communications networks, various types of signal processing, and autonomous weapons control are likely applications.

Factories are increasingly using programmable automation in production and management. For these systems to operate in an integrated manner, computers must be used to coordinate a myriad of tasks -- machine control, quality monitoring, scheduling, and problem diagnosis and correction -- all in time to affect the ongoing manufacturing process. The supercomputer's ability to perform complex operations in real-time makes it an attractive, if distant, possibility for automation.

Another possible use for supercomputers in this category is for space vehicles and planetary exploration. Researchers are working to make spacecraft "smarter", or more able to perform tasks without human presence or intervention. This requires vast quantities of computing power. For example, NASA is developing a Remote Orbital Servicing System; an unmanned spacecraft with robotic arms capable of performing maintenance on satellites. In its early stages, NASA expects to control this robotic manipulator device from Earth, but the long-term goal is to make the system capable of autonomous diagnosis and repair of satellites.

The academic research community has identified several areas of scientific endeavor that could be greatly enhanced by, or indeed, require supercomputers for research. Most of these areas of research are not now using supercomputers because only three U.S. universities, Colorado State, Purdue, and Minnesota, have them. By contrast, five universities in Western Europe had supercomputers at the end of 1982. The OTA Workshop participants reported several additional university purchases in both Western Europe and Japan.

Computer scientists require access to the most advanced examples of supercomputing equipment for the creation of languages, software, new architectures, and, indeed, new generations of supercomputers. Research scientists in high energy physics, molecular biology, chemistry, oceanography, meteorology, and other disciplines need access to at least today's machines for the furtherance of their disciplines. One of the most exciting aspects of a powerful new research tool is the potential for unexpected discoveries. It is actually quite likely that university needs will grow as more supercomputing power is made available to them because only hands-on experience will enable a researcher to understand what kinds of problems and which areas of research can be assisted with its use.

Currently, university access to supercomputers is limited by the machines' high price tags; between $5 and $15 million dollars. These prices do not include the applications software and the vast investments in personnel training and system experimentation required for effective user operation. The situation is not at all like a microcomputer in which a user need spend only a few hours in learning how to become facile in the machine's use. Effective user operation of a supercomputer can require many months of training.

Access Problems and Institutional Arrangements for the Provision of Access

Universities and smaller members of the industrial and Federal communities are all interested in access to supercomputer capabilities through networking. Commercial time-sharing arrangements for supercomputers have emerged; Boeing Computing Services, Control Data Corporation, and United Information Services have entered this market. The National Science Foundation is working on an arrangement for Federally-funded research. Recently, NSF formed a working group to study, among other topics, the possibility of establishing a national network for university research. The Department of Energy runs the Magnetic Fusion Energy Network based at the Lawrence Livermore National Laboratory. This network currently provides access to two Cray computers for energy research contractors; the network is expected to be expanded in the future.

Time-sharing on a supercomputer, if run like conventional arrangements, will diminish the unique problem-solving characteristics of the machine's architecture and logic. Supercomputers are designed to handle large, specialized numerical problems with particular types of regular structures such as the rectangular arrays of quantities involved in matrix algebra. Problems without these characteristics could be run on supercomputers, but there is no special advantage; in fact, there may be significant inefficiencies. Time-sharing arrangements that merely parcel out pieces of supercomputer power into small chunks will not realize the efficiences presented by the specialized architectures of these machines. In addition, the off-premise, dial-up nature of most computer time-sharing arrangements will inhibit (and may even prevent) user understanding of the machine's capabilities. It is this understanding, usually obtained through hands-on experience, that is thought to be critical for real gains in the advancement of those portions of science and engineering amenable to supercomputer assistance. Further, because supercomputers are not general purpose computing machines, arrangements must be organized around disciplines and/or common types of problem solving.

Supercomputer manufacturers have a great stake in university access to their machines aside from the revenues that might be generated by university purchase. University researchers using today's machines will test their architectural limits and provide manufacturers with information for new designs that can only be gained through actual experience.

The Adequacy of Research and Development in Supercomputers

In assessing the adequacy of research and development in supercomputers, it is useful to review the state of affairs in several contexts. Of course, all three contributors in this area, industry, universities and government, warrant examination. It is also important to separate comments about the adequacy of supercomputer R&D for the next generation from those about generations beyond. Finally, research and development of the next generation of supercomputers will not fulfill researchers' and industry's expectations without a concomitant software effort. Currently, application software development is taking place in each of the three communities. Access to the hardware will improve the university component of this effort, but a great deal more attention could usefully be paid to this important research area.

There are three U.S. producers of supercomputers: Control Data Corporation, Cray Research, Inc. and Denelcor, Inc. If one includes other specialized architectures such as array processors in the definition of supercomputers, there are about 40 additional producers. Producers of array processors include major corporations such as IBM and Control Data, as well as a group of small vendors like Floating Point Systems, Inc. The only other producers of supercomputers are in Japan; Fujitsu and Hitachi recently announced products for delivery next month, and NEC is said to be aiming for introduction in 1985.

The U.S. industry is small. This is undoubtedly a reflection of the present size of the market; at the end of 1982, there were only 60 supercomputers installed worldwide. Yet, there are some encouraging signs of health within the industry. Denelcor very recently entered the competition and has just sold its first machines. Since OTA's workshop, Control Data spun off ETA Systems, Inc., to produce the next generation of Cybers. Cray Research sold 16 machines in 1982, its 1983 sales are expected to number 30. The company originally expected to sell 25 Cray 1 machines; it has already shipped 55.

Development of the next generation machine is underway now in industry. But, serious concern about the ability of U.S. industry to keep its lead in the next generation has been raised. The Japanese National Superspeed Computer Project, sponsored by the Japanese government, is generally the source of this concern. The project's goal is to increase the theoretical speed of supercomputers by 1000 times by 1990. The theoretical speed of supercomputer machinery is not a particularly helpful benchmark against which U.S. efforts can be compared. The actual speed of a supercomputer is heavily dependent upon the applications software. Most researchers tell us that they are pleased to achieve between 20 and 50 percent of a machine's potential. There are, however, several other aspects of the U.S. supercomputer industry that can be examined for an understanding of the potential of U.S. performance:

(1) The United States certainly has a significant head start. The importance of user experience for design improvements and for the development of software, which is critical to the machine's overall performance, cannot be overemphasized. A reflection of the importance of this experience may be seen in the recent decision of Nippon Telephone and Telegraph to purchase U.S.-made supercomputers for delivery in 1985 despite the expected availability of a machine from NEC with nominally better specifications.

One of the on-going stimulants that enabled and sustains this lead has been the mission needs of the Department of Defense and the Department of Energy, a stimulant not found in Japan. DARPA has been funding advanced computer science research since about 1978, and expects to spend $750 million over the next ten years on various aspects of numerical and symbolic computer design. While this research focuses on artificial intelligence applications, there is a large area of commonality with supercomputer designs and applications, particularly in device technologies. A large portion of this research will take place in our universities and will be unclassified.

The Department of Energy has played a major role in supercomputer research through its design of nuclear weapons and controlled nuclear fusion research. These areas have added and will add a great deal to the understanding of supercomputer applications for many other areas of scientific endeavor, particularly those areas involving complex forms of simulation.

(2) Supercomputer development takes place among a group of almost uniquely American enterprises. The lean, competitive, sometimes flamboyant, environments of small, high-tech businesses seem to be of particular advantage in this kind of research. Indeed, the small business atmosphere is considered to be of such value that Control Data created such an entity (ETA Systems, Inc.) when the parent company felt it could not provide this environment through the larger structure. In addition, as is the case in many areas of high-tech, competition in this industry is very keen; not only from the three current actors, but from the pressure of potential entry by IBM, Trilogy and others.

(3) Informal, but very effective cooperation among the industrial, university and Government communities has emerged in supercomputer research. In support of its missions, the Federal Government has played major roles in supercomputer development as both purchaser of machines, and funder of research.

The purchasing role of the Federal Government has been important to the development of supercomputers. The very first available Cray 1 was installed at the Los Alamos National Laboratory without benefit of the attendant software. Since then, the government has purchased or has plans to purchase machines for weapons design, aerodynamic simulation, fusion energy research, propulsion system design, satellite weather and image processing, and a variety of other research questions.

This market opportunity provided by the Federal Government for equipment that is so costly and of such technical sophistication so as to otherwise have a limited commercial outlet, has obviously been and will continue to be important to U.S. industry. But the Federal Government's use of supercomputers in the Federal Laboratory System may have been of even greater value to industry. Because the machines were purchased as much with experimentation in mind as specific applications, much useful technological collaboration between purchaser and developer often ensued.

The Federal Government has funded research in supercomputers within both the National Laboratory and university communities. This research has proved to be of great value to industry, if in no other way than to have provided trained users and researchers. The National Laboratories and universities have been and are expected to be the environments in which new generations of supercomputers are developed. The vector processing architecture embodied in most of today's machines is only one of a variety of architectures now being considered for tomorrow's machines. There are many scientific and engineering problems with less regular mathematical structures that today's supercomputers cannot solve. The rigors of research and development under the particularly tight constraints of profit-making in this area do not generally allow industrial scientists to attempt more than incremental improvements in machine architecture.

For the development of new machine architectures, it is generally agreed that prototype machines must be built. In addition, prototypes are needed for software development. Estimates for the cost of a prototype project have been given as $7-$10 million to be spent over about a five year period. We know of no university prototype project with such funding. However, the number of on-going university projects in supercomputers is impressive; there are over 50 university-based projects. Most of these have very small amounts of funding, but are nevertheless training students in the design and potential of these complicated machines. A handful of schools have been funded at levels of more than $1 million per year.

In 1982, the Department of Defense provided about $8 million to five universities; The University of California at Berkeley, Stanford University, Carnegie-Mellon University, The California Institute of Technology and The Massachusetts Institute of Technology. Likewise in FY 1983, the Department of Energy provided about $3 million to six universities: New York University, The California Institute of Technology, The Massachusetts Institute of Technology, and the Universities of Illinois, Texas, and Wisconsin. While DARPA's and DOE's funding programs for future supercomputer research are currently being developed, they are expected to provide significant additions to these on-going university projects.

Future government participation, in both the funding of supercomputer research and in the purchasing of machines will provide a considerable hedge against decline in U.S. leadership in this area of technology. Government support of university research, particularly in the provision of access to supercomputers, is not only likely to have long-term positive consequences for our industrial competitiveness in supercomputers, but is also likely to assist in the advancement of scientific knowledge in a large variety of disciplines.

STATEMENT

OF

DR. PAUL B. SCHNECK

DIVISION LEADER

INFORMATION SCIENCES

OFFICE OF NAVAL RESEARCH

REPRESENTING THE U.S. ACTIVITIES BOARD

OF THE INSTITUTE OF

ELECTRICAL AND ELECTRONICS ENGINEERS

BEFORE THE

COMMITTEE ON SCIENCE AND TECHNOLOGY

U.S. HOUSE OF REPRESENTATIVES

NOVEMBER ?, 1983

Mr. Chairman and members of the Committee:

I am Paul B. Schneck, Division Leader, Information Sciences, of the Office of Naval Research, representing the Scientific Supercomputer Committee of the United States Activities Board, a major board of the Institute of Electrical and Electronics Engineers with access to IEEE's extensive resources and technical talents.

I am pleased to be here today at your invitation to discuss a report* prepared for USAB by the Scientific Supercomputer Committee; science and engineering problems requiring supercomputers; and mechanisms for increasing access to supercomputers for scientists and engineers.

The IEEE is the world's largest technical professional society with over 230,000 members worldwide, 190,000 of whom live and work in the United States. Thomas A. Edison was a founder of our organization, and our current membership has expertise in communications and computers as well as many other disciplines that have enabled the U.S. to land a man on the moon, create new industries, and achieve one of the highest standards of living in the world.

---

* SCIENTIFIC SUPERCOMPUTER COMMITTEE REPORT, Oct. 25, 1983, edition (See Appendix)

The IEEE USAB Scientific Supercomputer Committee Report

Early this year the United States Activities Board of the IEEE organized the Scientific Supercomputer Committee, chaired by Dr. Sidney Fernbach, consisting of thirteen individuals from industry, academia, and government, with expertise in all aspects of supercomputers including components, architecture, software, systems and peripherals, and use.

The Committee found that the economic and national security importance of supercomputers far exceeds their market share. The current direct applications of supercomputers to weapons design, weapons effects studies, and intelligence are critical to our strategic posture. Major emerging application of supercomputers to aircraft and automobile design, bio-technology, petroleum engineering, and semiconductor design are becoming fundamental to American industrial competitiveness. Further, each generation of mainline computer systems has descended from the prior generation of supercomputers. Advances in software and algorithms have similarly been carried forth onto mainline computer systems. To a great extent our Nation's leadership in the world computer marketplace, and success in that marketplace are dependent upon our supercomputer preeminence.

Supercomputer Applications Keyed to American Science & Engineering

Applications of supercomputers also contribute to our nation's economic and security postures. For example, NASA plans the Numerical Aerodynamics Simulation Program to improve the design and development of military and commercial aircraft, as well as space vehicles. Weather forecasting is done on supercomputers within the Department of Defense, NASA, and NOAA.

The Committee noted that supercomputers are beginning to replace physical experiments with numerical experiments. Computer simulation is often faster, cheaper, inherently safe, more flexible, and offers more insight than the methods which it is replacing. These factors have led to higher levels of design optimization, faster development cycles, lower use of materials and energy, and increased productivity. I will review later some examples of science and engineering requiring supercomputers cited by the Committee.

The Committee identified four technical areas that are impeding supercomputer development: architecture, peripherals, components, and software.

Today's supercomputers represent a major departure from past models. Previous systems had been able to take advantage of improvements in component technology. Vacuum tubes were replaced by transistors and transistors by (what are today called small scale) integrated circuits. Each change of technology considerably improved component speed. Components used in today's supercomputers operate at speeds measuring less than a billionth of a second.

Architecture

Two factors influenced computer architects to seek a major change. First, the pace of improving component speeds has slowed down. Second, even with current components, computers are limited by the speed of electricity - about ten inches in a billionth of a second. With traditional architecture, computers would be limited in speed even with arbitrarily fast components.

Computer architects first made use of techniques that were unseen by users. These techniques provided some extra hardware to increase the concurrency of operations within a computer to achieve higher speed. Further speed increases have demanded still more concurrency. Today's supercomputers are vector processors. Data are processed in an assembly line fashion, just as automobiles are manufactured. While very effective, this approach imposes programming constraints on the user and cannot be scaled up for further improvement. The next step in increasing the speed of supercomputers is to link them together as multiprocessors. This is just beginning to happen at the present time; one supercomputer manufacturer (Cray Research) has recently delivered its first multiprocessor model. While it is hoped that this approach can be scaled up significantly there is very little experience in algorithms of programming for multiprocessors.

Research in advanced architectures is limited by substantial lag in software development, poor access to current supercomputers, and the high cost associated with prototype construction. In the case of both peripherals and components, the supercomputer market is too small to overcome the market forces that pressure suppliers to move in other directions.

Peripherals

Peripherals are used to store data and on demand transfer data rapidly to or from the computer. The speed and capacity of peripherals have not kept pace with the increased capability of supercomputers. Already, some problems are limited by the performance of current peripherals. Limitations will be even more severe in the future as the imbalance continues to increase.

Components

Components are another area where supercomputer manufacturers have unsatisfied requirements. Because of the small numbers of parts required and their high speeds and manufacturing complexity, the merchant semiconductor companies are not devoting enough resources to meet the needs of the supercomputer manufacturers. They are concentrating on production of the memories and microprocessors used in personal computers and other consumer electronics. Dependence on foreign components would leave U.S. manufacturers vulnerable to external political issues, and purchasing custom components is already resulting in transfer of proprietary information overseas. Cray Research, Inc. has established its own semiconductor operations to obviate this problem, at the expense of its other important research projects.

Software

The first vector processor supercomputer was installed seven years ago. Software and applications have been slowly improving since that time. The availability of nearly 100 vector supercomputers allows more opportunity for experimentation and learning. The pace has been slow. Ready university access to vector processors could greatly speed up the development of advanced software and applications.

Multiprocessor supercomputers will be appearing throughout the rest of this decade. The issues we face in software and algorithms are even more difficult. In many cases we have not yet determined the direction we should take. Many experiments will be necessary. Access to multiprocessor supercomputers at the universities will hasten the development of software and applications.

Foreign Competitors

The Committee examined foreign competition in supercomputers and supporting areas of information technology. Several countries have embarked on projects to design and build supercomputers. France, through the Ministry of Defense, has projects exploring vector processors, multiprocessors, and their combination. Three Japanese companies have announced plans for vector supercomputers. The peak speed of each will rival or exceed that of current supercomputers, while their planned memory capacity is eight times as great. The United Kingdom plans to pursue "A Programme for Advanced Information Technology" including specific efforts in components. Finally, the Commission of the European Communities has embarked on a joint program, The European Strategic Program for Research in Information Technology, pooling their resources and minimizing duplication.

Science and Engineering Problems Requiring Supercomputers

Some engineering and science problems which are highly regular in their mathmatical formulation are already being addressed with supercomputers or are likely candidates. Such problems in science include weather forecasting, molecular chemistry, genetics, mathematics, and economics. In engineering such problems include VLSI design, structural analysis, seismic analysis and reservoir modeling, transmission line design, and image processing for robotics, flight simulators, and the motion picture industry.

Foremost among the problems to which supercomputers are being applied is nuclear weapons design. This problem features high return, not attainable with conventional computing equipment; regular structure, permitting effective use of supercomputers; and advantages in simulation over testing. Similar qualities are seen in studies of weapon effects, magnetic fusion energy research, and nuclear reactor design and development.

Computational fluid dynamics applied to the design of aircraft and space vehicles is recognized as benefitting greatly from supercomputer availability. Wind tunnels, the traditional test mechanism, are expensive to build and operate, suffer from the inability to modify rapidly a model to incorporate a new idea, and provide only global information such as total drag, not local information such as drag components. Supercomputer power allowed the Boeing Company to isolate "interference drag" and identify its source, permitting a close coupled nacelle position on the reengined 737-300. This has allowed smaller, lighter landing gear, and lower drag resulting in significant production and operating economies. Supercomputer power also permitted design of the 757 cab compatible with the 767 cab in the short time frame available. This resulted in a common type-certificate, permitting a pilot certified on one aircraft to operate the other, and yielding a significant competitive advantage. For every dollar spent on wind tunnel testing Boeing spent twenty-five to thirty cents on computation. That difference will continue to narrow.

Twenty years ago Richard Hamming noted "The purpose of computing is insight, not numbers." Supercomputers may offer the biggest advantages in areas in which, uncertain of the detailed science, we create empirical models. Just as access to the supercomputer's capability allowed Boeing scientists and engineers to experiment and understand the cause of interference drag, access to supercomputers will allow other researchers to experiment to test their understanding. By making an experiment's time scale small enough to be consistent with human time scales, a supercomputer facilitates discovery.

Access an Important Issue for Scientists and Engineers

The Committee noted that one of the most important issues in supercomputing today is the need for access. The United States is educating and training a generation of scientists and engineers with no experience in the use of supercomputers. Twenty years ago supercomputers, of their day, were being used on many American campuses. Today there are only three. We are not preparing the science and engineering community to work with the most advanced machines that are available. At many institutions small processors are shared among students and faculty. The power available to an individual is severely limited, but even more importantly there is no opportunty to gain experience with the special characteristics of today's supercomputers.

As previously mentioned, the architecture of supercomputer software and hardware is very different from ordinary computers. Because of the high cost of supercomputers, only a small number of engineers and scientists have the opportunity now to learn how to use them. Financial support from the U.S. Government for the establishment of a number of supercomputer centers should be provided. These centers typically would be available to universities, which would provide student and faculty access. Engineers, physical and social scientists, mathematicians, and people from all disciplines would learn to simulate complex problems using the supercomputers. U.S. industry would benefit. The Committee believes that much broader use of supercomputers is essential and can be accomplished only if more trained people are available. Establishment of several centers would be advantageous to encourage diversity and foster more extensive interaction with the center staff. Sharing of software would be encouraged by the common use of central facilities.

The issue of increased access in the industrial sector is more complex. It is a reflection of the scarcity of personnel with training or experience on supercomputers, as well as the common need to cost justify the acquisition of a supercomputer for a specific program. There is a tremendous base of existing applications software; much software development consists of modifying or augmenting that base. Programs written for conventional computers often require significant overhaul before they are cost effective on a supercomputer. With the pressures of time and limited personnel resources it is frequently not economic to convert a program for use on a supercomputer. In such cases the introduction of supercomputers will be delayed until trained personnel are available, software is developed to automatically convert programs, supercomputers become more cost effective with software for conventional machines, or problem criticality increases dramatically. The government role in this area is indirect. By supporting the development of all elements of supercomputer systems (components, architecture, software and applications, and peripherals) progress towards cost effectiveness for existing software is accelerated.

There is another kind of access that must be addressed. New architectures with unfamiliar names and with unfamiliar software needs are being studied at a number of universities. When these machines are built they should be attached to a network, e.g. ARPANET, allowing access to the research community. In this way different approaches can receive scrutiny and feedback from a heterogeneous population of engineers and scientists. Federal agencies with strong mission requirements for computing should be supported in acquiring such experimental systems. If promising, the entrepreneurial marketplace will continue their delelopment and commercialization.

Government Role and Recommendations

The Committee found that the national importance of supercomputers exceeds that indicated by the size of their market. Development of supercomputers must compete for the same talents used in the creation of consumer electronics or mainline commercial products and cannot promise the same scale of profits or time. Few companies have been motivated to enter this marketplace; most remain active only in the mainline market. To assure a continuing development cycle for all aspects of supercomputers, the Federal government must provide sufficient incentive, in keeping with the national importance of supercomputers.

The Committee concludes its report with five recommendations:

1. That the Federal Government make a long range commitment to maintain leadership in scientific supercomputer development. This point addresses the economic and national security importance of supercomputers and recognizes that to protect its security the U.S. must be second to none in the computing resources it applies to its interests.

2. That the Federal Government encourage both the continued development of scientific supercomputers and their availability to industry through the research community. Therefore, the Government should purchase (through its agencies) a minimum of 12 highly advanced scientific supercomputer systems specified to have significantly increased performance. The mechanisms available to the government include tax incentives, direct support, and clarification of laws relating to joint research ventures. The stimulation of supercomputer production coupled with increased access to the research and academic communities will encourage further growth in the marketplace.

3. That the Federal Government support the development of all elements of supercomputers systems. Progress in the areas of components, architecture, software and applications, and peripheral equipment is necessary to maintain broad applicability of supercomputers and to permit their continued evolution and development.

4. That the Federal Government designate a lead organization to coordinate the roles of the many Federal agencies which are critically dependent on supercomputer systems. The Department of Energy has a long and successful history of bringing new supercomputers into productive use.

5. That the Federal Government establish continuing national goals for the development of scientific supercomputer systems well into the 21st century. In the face of foreign competition a strong, well defined program of national goals is necessary to assure that all facets of supercomputing receive adequate attention and resources. These goals should include the setting of state-of-the-art supercomputer specifications to be met at the end of each five-year period during that time. The development of efficient software techniques for utilizing the hardware architecture should be specified at the same time. Purchase by the government of these advanced systems should be assured at the time the specifications are accepted by the manufacturer.

Conclusion

Mr. Chairman, thank you for this opportunity to present the IEEE United States Activities Board's views on these issues of national importance. I would be pleased to answer any questions at this time.

SHORT BIOGRAPHY

Alvy Ray Smith
Computer Graphics Project Leader

Lucasfilm Ltd.

Alvy was introduced to computer graphics at Xerox Palo Alto Research Center by Dick Shoup, who designed the first sophisticated computer-assisted "paint" machine (1972-1973). Here he united his artistic interests and computer science education (PhD, Stanford, Cellular Automata Theory) to create the videotape *Vidbits* which served as a passport into the newly formed Computer Graphics Laboratory at the New York Institute of Technology in 1975, where for five years he assisted in the development of this facility into one of the premier computer graphics installations in the world. While there he generated stills, videotapes, and films which have appeared in a variety of magazines, books, newspapers, museums, and televsion shows. In particular, just before leaving NYIT, he collaborated with Ed Emshwiller on the videotape *Sunstone*. He joined Ed Catmull in the Computer Division of Lucasfilm Ltd. in early 1980. This division is involved in a major effort to digitize the filmmaking process with special emphasis, of course, on the use of computer graphics in theatrical release motion pictures. Alvy is Project Leader of the division's Computer Graphics Project, consisting of thirteen people engaged in both software and hardware development. The Graphics Project generated a state-of-the-art production, under Alvy's direction, of a sequence (the "Genesis Demo") for the 1982 Paramount movie *Star Trek II - The Wrath of Khan*. Most recently, the Project contributed a sequence to Lucasfilm's *Return of the Jedi*.

COMPUTER POWER FOR FILM AND FLIGHT [†]
Technical Memo No. 95

Alvy Ray Smith
Computer Graphics Project Leader

Computer Division
Lucasfilm Ltd.
October 21, 1983

INTRODUCTION

Recent incorporation of computer graphics into feature-length films - *Return of the Jedi, Star Trek II: The Wrath of Khan,* and *Tron,* to name a few - has inspired a belief that entire films might soon be generated by computer. It is shown below that even the so-called "supercomputers" of today fall quite short of the power required for this goal. True supercomputers are needed with capabilities just now being conceived and with cost commensurate with typical filmmaking practice.

The discussion here will focus on the computation of images, but supercomputers are general-purpose computers, which means they can be used for any kind of computation. In particular, supercomputers meeting the demands of movie-making computation would further image-making computation in general. Flight simulators are a case. Cost-effective supercomputers of sufficient power could replace expensive special-purpose flight simulators.

HOW TO COUNT PIXELS

A *pixel* is the smallest part of a picture computed by a computer (it is a corruption of "picture element"). A single frame of film is assumed to be divided into a great many scanlines - just as a television screen is divided into scanlines - and each of these is further divided into many tiny spots, the pixels. So a frame is divided into, say, 1000 scanlines each consisting of, say, 1000 pixels.

The *resolution* in this case would be said to be 1000×1000 ("×" is read "by"), or 1K×1K, where K is shorthand for a thousand of anything [‡]. The resolution required by film is not yet known with great certainty. Since 500×500 is about the resolution of American television, this is undoubtedly too low for film. 8K×8K is surely too high. We shall assume it is 3K×1.5K pixels. The Japanese have experimented with resolution in a series of tests for establishing their *hi-vi*, or high-definition television system. They found that above 1600 scanlines their subjects did not notice substantial improvements in image quality. For ease of computations and conservatism, we will use 1500 scanlines for our estimates here. Since a common movie film format is Panavision, or Cinemascope, which has a width more than twice its height, we shall assume 3000 pixels along each scanline.

The number of pixels per frame is simply the number of scanlines times the number of pixels on a scanline. So a typical frame is assumed to have 4.5 million pixels; a computer would have to compute 4.5 million spots to make such a frame.

[†] Prepared for testimony before the House of Representatives Committee on Science and Technology, Congressman Don Fuqua Chairman, Washington, D.C., October 26, 1983.

[‡] In computer science, K usually abbreviates 1024, the tenth power of 2; we use K = 1000 throughout this article, however.

Another abbreviation we will use is 1M for one million; so 4.5M pixels is 4.5 million pixels. Another common way to say the same thing is "4.5 megapixels". So M and mega- mean one million. There are similar abbreviations for larger and larger numbers. Since we will use very large numbers in this article, it will help to abbreviate them as shown in the following table, where each succeeding entry is three orders of magnitude larger than the one above (a factor of 10 is one order of magnitude):

Large Number Abbreviations			
Abbreviation	Prefix	Meaning	No. of 0's
K	kilo-	thousand	3
M	mega-	million	6
G	giga-	billion	9
T	tera-	trillion	12
P	peta-	quadrillion	15

HOW MANY BYTES IN A MOVIE?

A color film has a red, green, and blue primary component at each pixel to give full color. To "compute a frame" means that a computer must compute one red, one green, and one blue item for every pixel in a frame. The typical smallest unit a computer can generate is called a *byte* (eight bits). So at the very minimum, computing a frame requires the generation of 3 bytes at every pixel, or 13.5M bytes for an entire frame.

It is fairly well known that film is projected at 24 frames per second; so it is not difficult to calculate that there are 24 times 60 times 90, or 129,600, frames in a typical 90-minute movie. Using the abbreviations, there are 129.6K frames in a typical movie. So computing a movie requires computing 13.5M times 129.6K bytes, or 1.7496T bytes. *A typical 90-minute movie requires the computation of about 1.75 terabytes.* This very large number is the amount of memory storage space required to store a movie digitally. Next we estimate the computer power necessary to compute this many bytes of finished film.

HOW LONG WOULD IT TAKE TO COMPUTE A MOVIE?

The calculation above shows that a typical feature-length film at a reasonable resolution has over a trillion bytes. This is the size of the finished film; mistakes, reshoots, and cutting-room floor footage are not counted. Each byte of finished film requires the execution of a series of operations by a computer. The operations of a computer are controlled by a list of instructions, its *program*. Thus the speed of a computer is sometimes measured in *instructions per second*. This works for today's slower general-purpose computers which tend to perform one operation per instruction, but supercomputer speeds are measured in *operations per second* since they tend to to perform many operations for each program instruction.

We estimate that a computer must execute from 1000 to 10,000 operations for each final output byte to film. This is based on an image of a richness we believe filmgoers will find satisfactory, one which approaches reality in complexity (but certainly not restricted to reality in content). We shall use the average of 5000 operations per byte for our calculations. This is believed to be an underestimate since we used about 4000 operations per byte for the Genesis Demo sequence in *Star Trek II: The Wrath of Khan* and are aiming for an even higher complexity in future projects. 5K operations per byte times about 1.75T operations implies 8.75P operations, or 8.75 peta-operations. Thus *a typical movie would require the computation of approximately 8.75 quadrillion operations.*

A common general-purpose computer in 1983 (not a supercomputer) is the VAX 11/780 manufactured by Digital Equipment Corporation. It is rated as a 1Mip machine. This is pronounced "one mip" and means 1M instructions per second. For a VAX, 1 Mip is the same as one million operations per second; so 8.75P operations would take 8.75G seconds to compute. Since

$$1\text{G second} = 11.574\text{K days} = 31.797 \text{ years},$$

a VAX would require more than 278 years to compute a movie! And this is being kind. As we will point out several times in this memo, it is inaccurate to use the rated speed of a computer for realistic results. The VAX computes at much less than 1 Mips for actual programs; so 278 years is an unrealizable best case.

Since it is out of the question to compute an entire 3K×1.5K movie at the present, what if video resolution, 500×500, were used instead? A best-case calculation similar to that above yields the result that over 15 years of VAX time would be required; so even a highly degraded image is also out of the question.

It should be emphasized here that we are calculating *raw compute time*. We have assumed that the machine never stops computing once it starts. This is unrealistic because a computer never operates around the clock, day after day, without having a hardware failure or being stopped for some other reason, such as machine maintenance. Furthermore, time must be allotted for system security backups, recomputations for the inevitable mistakes, and program development time.

A typical major motion picture requires three stages called preproduction, production, and postproduction. Preproduction comprises such tasks as design, scripting, storyboarding, and financing; it takes on the order of six months. Postproduction covers editing, audio mixing, film printing, foreign language adaptation, special effects, advertising, and distribution; it requires on the order of a year. Our raw compute time takes neither of these two stages into account; it is to be compared to the actual production time of a film which is about four to six months.

THE EFFECT OF SUPERCOMPUTERS

The numbers derived above are discouraging. Would the use of modern supercomputers change the outlook substantially? The two most well-known supercomputers today are the Cyber 205 by Control Data Corporation and the Cray 1S by Cray Research Inc. The speeds of these machines - and of supercomputers in general - are expressed in "megaflops" instead of, say, "mega-ops" because the operations performed by supercomputers which most interest their users are the so-called "floating-point operations", a particularly difficult set of operations for computers. A supercomputer megaflop is computationally more powerful than a VAX Mip, but the difference is difficult to quantify; so we will take them as equal here. The effect is to underestimate the time needed to compute with a VAX rather than to overestimate the time needed with a supercomputer.

A recent paper [1] argues that in actual practice at Los Alamos National Laboratory, the speeds of these two machines are more like 10-20 megaflops. We shall use 50 megaflops here in an attempt to use a number both benign and realistic. It is important to note that this is substantially lower than the advertised speeds. For example, both American supercomputers have advertised speeds of several hundred megaflops. Parts of some programs can indeed compute this rapidly on them but real programs cannot effectively sustain this speed. As reported in [1], both machines run at less than 20 megaflops on a real mix of programs.

With the assumptions made above, a Cray or Cyber is about 50 times more powerful than a VAX. Thus, *a modern "supercomputer" would take about 6 years to compute a feature-length film*. This is still entirely too long.

WHAT ABOUT JAPANESE SUPERCOMPUTERS?

The Japanese firms of Fujitsu and Hitachi plan to have the supercomputers FACOM VP-200 at 500 megaflops and HITAC S-810 at 630 megaflops, respectively, operational by the end of 1983 [2]. If the top speeds are assumed, our hypothetical movie would take about 29 or 23 weeks to compute, respectively, but, as argued above, top speeds are not to be trusted for realistic speed measurements of supercomputers. In actuality, these machines will probably be comparable in speed with the Cray and Cyber supercomputers and hence too slow for full-length movies.

The Japanese have announced their intention of building a 10 gigaflop computer by 1989. The Super-Speed Scientific Computer Project has a total budget of about $100 million [2]. A 10G flop machine is about 20 times as powerful as a current supercomputer. The raw compute time for a film thus reduces to about 100 days, which is approaching a reasonable number for a feature film. *A supercomputer of the power planned by the Japanese would be sufficiently powerful to compute a feature-length film.* Both Cray Research and ETA (the recent supercomputer spinoff company of Control Data) have announced intentions of producing supercomputers of similar power by the late 1980s or 1990[3]. ETA, in fact, is planning a 30 gigaflop machine.

THE COST OF COMPUTING MOVIES

A conventional movie costs on the order of $10M with about $8M for production and $2M for postproduction and preproduction. Each one takes about two years. Only one of every ten movies succeeds on the average, but we shall not incorporate this unfortunate fact into our calculations here.

A current supercomputer installation costs $1-2M for initial site preparation and peripheral equipment plus about $3-4M per year for staff, building lease, supplies, utilities, and supercomputer maintenance. The machine itself costs about $10M; current prices range from $4M to $12M. Spreading the one-time costs of about $12M over three years means a per year cost of $4M. Thus the annual cost is $7-8M for a current supercomputer facility.

Comparing the production costs of a typical film with this annual supercomputer cost shows them approximately equal, if the machine is used for computing a new (successful) film every year. Of course, as we have seen above, current supercomputers cannot compute a film in a year. *So what we are awaiting are supercomputers of the power planned for 1990 at current "supercomputer" prices or less.*

FILM VERSUS FLIGHT

The production of complex graphics for films and realistic graphics for flight simulators - simulators for commercial aircraft, military aircraft, space shuttles, and even oil tankers - obviously have much in common. There are major differences, however. In film, a single path is taken through a rich terrain where the terrain changes from film to film; the time taken to compute a frame is of secondary importance to the richness of the image, several minutes per frame being tolerable. In flight, multiple paths are flown through a terrain which does not change from flight to flight; the time per frame is strictly that of *real time,* 24 frames per second for film or 30 frames per second for video, and image complexity is sacrificed to meet this schedule. Film implies great complexity at a cost of time, while flight implies real time at a cost of complexity. The data complexity for a film is concentrated along the one predetermined path through the database; Hollywood is famous for its false fronts. Flight simulators must spread data complexity over the wide range of potential paths which are not known in advance (except to within restricted "corridors" through the data).

COMPUTATIONAL REQUIREMENTS OF A FLIGHT SIMULATION

If a flight simulator were to display at the resolution used above for films, 13.5M bytes per frame, and at the complexity used above, 5000 operations per byte, then it would have to compute at a rate of 2.025 teraflops to satisfy the demands of real time (30 frames per second). This is four orders of magnitude faster than current "supercomputers" and two orders of magnitude faster than the projected 10 gigaflop supercomputers. So it is not feasible to even consider the use of a general-purpose supercomputer, current or projected, for flight simulations at movie resolution and complexity.

As already acknowledged, flight simulators typically sacrifice image complexity and resolution for speed. A calculation similar to that above but at video resolution (750K bytes per frame) and one-tenth the complexity (500 operations per byte) yields a computation rate of 11.25 gigaflops, in the ballpark for a 10-30 gigaflop machine. Therefore, *it is conceivable that a supercomputer of the power projected for the 1990s could compute a flight simulation in real time.*

This is important. Current flight simulators are special-purpose machines, generally one of a kind or a few of a kind. We are suggesting that a general-purpose supercomputer could conceivably replace these in the coming decade or two. A change would be a programming change only, not a new hardware development requiring great cost in time and money.

CONCLUSIONS

We have argued that the goal of completely computing a feature-length film (90 minutes) at a reasonable resolution (3K by 1.5K pixels) in a reasonable amount of time (2-3 months, nonstop) is not feasible with today's "supercomputers" but would be with supercomputers of the power planned by the Japanese and Americans for 1990. This assumes that the cost of such a machine would be about $5 to $10 million. This means that, without a special boost, we are unlikely to realize this particular goal in the US until about 1995-2000.

We have further argued that any general-purpose approach that improves movie-making computations also improves all image-making computations. Flight simulation, a case in point, could conceivably be computed on general-purpose supercomputers of the variety planned for the 1990s, although at a resolution and complexity less than that for film.

The main point is that it is not at all difficult to imagine what we would do with the supercomputers apparently still a decade away. We in the image-making professions are preparing ourselves for cost-effective supercomputing; the hunger has begun.

REFERENCES

[1] "The Computational Speed of Supercomputers", Ingrid Y. Bucher, *SIGMETRICS 83 Conference Proceedings*, ACM, August, 1983, Minneapolis, pp. 151-165.

[2] "Overview of Scientific Computing System - Next Generation Japanese Supercomputers", Hiroshi Kashiwagi, *Digest of Papers: Spring COMPCON 83*, IEEE, February 28-March 1, 1983, San Francisco, p. 305.

[3] "The Battle of the Supercomputers: Japan's All-Out Challenge to the U.S.", *Business Week*, October 17, 1983, pp. 156-166.

PROPOSED TESTIMONY OF DR. GEORGE A. KEYWORTH, II
SCIENCE ADVISOR TO THE PRESIDENT, AND
DIRECTOR, OFFICE OF SCIENCE AND TECHNOLOGY POLICY
EXECUTIVE OFFICE OF THE PRESIDENT

TO THE COMMITTEE ON SCIENCE AND TECHNOLOGY
UNITED STATES HOUSE OF REPRESENTATIVES

HEARINGS ON SUPERCOMPUTERS
NOVEMBER 16, 1983

CHAIRMAN FUQUA AND MEMBERS OF THE COMMITTEE:

I'M PLEASED TO HAVE THIS OPPORTUNITY TO MEET WITH YOU AND TO OUTLINE THE ADMINISTRATION'S VIEWS AND PLANS CONCERNING SUPERCOMPUTERS. TODAY I WILL DISCUSS THE PRINCIPAL ELEMENTS OF THE ADMINISTRATION'S PROGRAM TO ENSURE CONTINUED U.S. TECHNICAL LEADERSHIP IN THIS AREA. I WILL ALSO REVIEW OUR STRATEGY FOR THE SEPARATE, BUT RELATED, QUESTION OF HOW TO PROVIDE ACCESS TO THESE INCREASINGLY IMPORTANT RESEARCH TOOLS TO MORE OF OUR TECHNICAL COMMUNITY. EARLY NEXT YEAR THE PRESIDENT'S FY 1985 BUDGET WILL PROVIDE THE DETAILS OF HOW THE ADMINISTRATION WILL IMPLEMENT THESE POLICIES. NEVERTHELESS, IN THE MEANTIME I BELIEVE THAT IT IS IMPORTANT AND USEFUL FOR ME TO LAY OUT THE FRAMEWORK WITHIN WHICH THE ADMINISTRATION IS ADDRESSING THIS ISSUE, AND TO IDENTIFY KEY POLICY EMPHASES THAT WILL COMPLEMENT THE VERY CONSIDERABLE CAPABILITIES AND RESOURCES OF THE PRIVATE SECTOR. I BELIEVE THAT WHAT WE ARE PROPOSING CONSTITUTES A LOGICAL AND REASONABLE NATIONAL PROGRAM, ONE THAT WHICH WILL ACCOMPLISH OUR GOALS OF CONTINUED TECHNICAL PRE-EMINENCE AND SCIENTIFIC LEADERSHIP.

SUPERCOMPUTERS HAVE TRADITIONALLY BEEN A TECHNOLOGY COMPLETELY DOMINATED BY U.S. INDUSTRY. AS SUCH, THEY SYMBOLIZE U.S. LEADERSHIP IN MANY OF THE HIGH TECHNOLOGY AREAS THAT ARE CENTRAL TO OUR FUTURE NATIONAL SECURITY AND INTERNATIONAL ECONOMIC COMPETITIVENESS. RECENTLY SUPERCOMPUTERS HAVE ATTRACTED HEIGHTENED ATTENTION BECAUSE OUR HISTORIC DOMINATION OF THE WORLD MARKET IS NOW BEING CHALLENGED. JAPAN HAS ANNOUNCED ITS INTENTION TO INTRODUCE OVER THE NEXT SEVERAL YEARS COMPUTERS WHOSE STATED PERFORMANCE OBJECTIVES ARE TO MATCH THE U.S. MACHINES. HOWEVER, THOSE ACTUAL CAPABILITIES HAVE YET TO BE DEMONSTRATED, AND IT WILL PROBABLY BE ANOTHER YEAR OR SO BEFORE THE PERFORMANCE NUMBERS WE HAVE SEEN PUBLISHED FROM TIME TO TIME CAN BE COMPARED TO THE KNOWN CAPABILITIES OF CURRENT U.S. PRODUCTS IN VARIOUS APPLICATIONS.

THE ADMINISTRATION RECOGNIZES THE IMPORTANCE, FOR THE UNITED STATES, OF MAINTAINING OUR LEAD IN DEVELOPING SUCCESSIVE GENERATIONS OF THESE EVER-FASTER, GENERAL PURPOSE COMPUTING MACHINES. IT INTENDS TO CONTINUE TO MAINTAIN A CLIMATE IN WHICH U.S. SUPERCOMPUTER COMPANIES WILL HAVE THE OPPORTUNITY AND BE ENCOURAGED TO REMAIN IN THE LEAD.

ABOUT A YEAR AGO, I ESTABLISHED AN AD HOC COMMITTEE OF THE FEDERAL COORDINATING COUNCIL ON SCIENCE, ENGINEERING AND

Technology (FCCSET) to look into the various concerns which were being voiced about supercomputers. I asked them to advise me on what the Government's role might be to help retain the U.S. competitive edge in this vital field of technology. That committee in turn established three interagency working groups, representing the agencies who have a significant interest in supercomputers: DOE, DOD, NSF, NASA, DOC and the intelligence community. We are today making publicly available, together with a brief cover statement, the reports of two of those working groups. The third group has the on-going task of coordinating supercomputer related R&D among the agencies, and no report has been requested. The FCCSET reports are just two of the many inputs that we have received and considered in developing our program. We have also consulted extensively with industry, including both current manufacturers of supercomputers and major firms in related fields, with the universities, with the key Federal laboratories, and with many knowledgeable individuals throughout the country.

In the forthcoming fiscal year 1985 budget, President Reagan will propose a variety of activities to enhance U.S. R&D strengths and to encourage U.S. supercomputer firms to continue to outpace their competitors. Those Federal programs will be a logical extension of ongoing activities related to supercomputers and will not propose government intervention into the area

OF INDUSTRY'S OWN CHOICE AND JUDGMENT. THE GOVERNMENT ROLE WILL BE GUIDED BY ITS OWN SPECIFIC NEEDS FOR SUPERCOMPUTERS AS THE WORLD'S LARGEST USER OF THESE SYSTEMS AND BY THE WELL-ESTABLISHED COMMITMENT TO SUPPORT THE WORLD'S LEADING SCIENCE AND TECHNOLOGY BASE.

THE ADMINISTRATION IS PROPOSING TO ADDRESS TWO FUNDAMENTAL PROBLEMS. ONE IS TO ENSURE THE CONTINUED AVAILABILITY TO THE GOVERNMENT OF U.S. MANUFACTURED SUPERCOMPUTERS ABLE TO SATISFY NATIONAL SECURITY AND OTHER CRITICAL RESEARCH NEEDS. MACHINES WITH APPROXIMATELY TWO ORDERS OF MAGNITUDE GREATER CAPABILITIES THAN PRESENTLY AVAILABLE ARE ALREADY BEING PLANNED FOR. THE SECOND GOAL IS TO PROVIDE BROADER ACCESS BY U.S. RESEARCHERS TO THE LATEST MACHINES, WHICH IN AN INCREASING NUMBER OF CASES ARE NECESSARY FOR PERFORMING STATE OF THE ART SCIENCE AND ENGINEERING RESEARCH.

HOW DO WE PLAN TO ACCOMPLISH THESE GOALS? THE APPROACH TO THE AVAILABILITY OF SUPERCOMPUTERS FOR DIRECT GOVERNMENT NEEDS HAS THREE MAJOR COMPONENTS, WHICH WILL BE ADDRESSED IN COMPLEMENTARY FASHION PRIMARILY IN THE BUDGET PROPOSALS OF THE DEPARTMENT OF ENERGY, THE NATIONAL SCIENCE FOUNDATION, AND THE DEPARTMENT OF DEFENSE. IN THE SHORT TERM THE GOVERNMENT'S OWN NEEDS FOR SUPERCOMPUTERS WILL REQUIRE THE CONTINUING ACQUISITION OF THE FASTEST MACHINES CURRENTLY AVAILABLE, AND AGENCIES WILL COORDINATE THEIR NEEDS FOR THESE

MACHINES. INDUSTRY HAS REPEATEDLY EMPHASIZED THAT GOVERNMENT CAN BEST SUPPORT CONTINUED ADVANCES BY ESTABLISHING CLEAR STATEMENTS OF ITS OWN REQUIREMENTS. IN THIS CONTEXT, THE ADMINISTRATION RECOGNIZES THE ADVANTAGE OF CLOSE RELATIONSHIPS BETWEEN INDUSTRY AND THE GOVERNMENT LABS AND UNIVERSITIES--A RELATIONSHIP THAT SERVED THE NATION SO WELL IN THE EARLY DAYS OF COMPUTING. PROTOTYPE OR PRE-PRODUCTION MACHINES CAN BE INSTALLED IN FEDERAL LABORATORIES OR UNIVERSITY CENTERS AND SOFTWARE CAN BE JOINTLY DEVELOPED FOR THESE NEW MACHINES. WHEN THE PRODUCTION VERSIONS OF THESE MACHINES ARE LATER INSTALLED TO MEET BROAD GOVERNMENT REQUIREMENTS, THEY DON'T THEN REQUIRE SEVERAL FRUSTRATING YEARS OF CODE DEVELOPMENT BEFORE THEY CAN CONTRIBUTE EFFECTIVELY.

IN THE LONGER TERM THE GOVERNMENT WILL SET PERFORMANCE GOALS FOR ITS FUTURE PURCHASES, POSSIBLY IN THE CONTEXT OF THE PROBLEMS IT NEEDS TO SOLVE. THIS APPROACH HAS THE ADVANTAGE OF SETTING CLEAR TARGETS WITHOUT TRYING TO PREJUDGE THE MOST EFFECTIVE WAYS OF REACHING THEM IN A FIELD THAT IS UNDERGOING A FUNDAMENTAL SHIFT IN TECHNOLOGY FROM SERIAL TO PARALLEL PROCESSING. IT ALSO PERMITS MANUFACTURERS TO RESPOND TO GOVERNMENT NEED WITH SYSTEMS THAT COULD SATISFY A DEMAND IN THE COMMERCIAL MARKETPLACE.

ONE OF THE FCCSET REPORTS RECOMMENDS GUARANTEED GOVERNMENT PURCHASES OF ALL SYSTEMS MEETING PRE-SET PERFORMANCE GOALS. HOW THIS PROPOSAL CAN BE EQUITABLY IMPLEMENTED NEEDS FURTHER STUDY. ONE APPROACH THAT HAS BEEN SUGGESTED IS FOR AGENCIES TO ESTABLISH EXPERIMENTAL PROGRAMS, WHERE THEY COULD DISCUSS SPECIFIC PROBLEMS AND SUPERCOMPUTER NEEDS WITH ONE OR MORE VENDORS. THEY WOULD THEN ENTER INTO SOME KIND OF CONTRACTUAL AGREEMENT WITH A VENDOR TO SUPPLY ONE OR MORE SYSTEMS, WITH MUTUALLY AGREED-UPON SPECIFICATIONS, DESIGNED TO FILL THE IDENTIFIED NEEDS. WHILE IN SUCH A CASE THE GOVERNMENT IS REDUCING THE UNCERTAINTY OF THE MARKETPLACE BY GUARANTEEING A BUY--OF ONE SYSTEM AT LEAST--THE VENDOR STILL HAS THE TECHNOLOGICAL RISK OF MEETING THE SPECIFICATIONS, FAILING WHICH, THE TAXPAYER WOULD HAVE NO OBLIGATION TO PAY. WE WILL BE EXPLORING THIS AND OTHER APPROACHES WITH THE AGENCIES.

THE CONTINUING EVOLUTION OF MACHINES WITH CAPABILITIES SEVERAL ORDERS OF MAGNITUDE GREATER THAN THOSE CURRENTLY AVAILABLE REQUIRES MARKET INCENTIVE AND CAPITAL FOR THE COMPUTER ENTREPRENEURS. BUT EQUALLY IMPORTANT, ONGOING PROGRESS WILL REQUIRE A CONTINUED SUPPLY OF OUTSTANDING COMPUTER SCIENTISTS, AND IT WILL REQUIRE CONTINUED SUPPORT OF FIRST CLASS RESEARCH IN COMPUTER SCIENCE AND ENGINEERING. IN BOTH THESE AREAS THE FEDERAL GOVERNMENT HAS A MAJOR RESPONSIBILITY, AND THE ADMINISTRATION EXPECTS TO PROPOSE SEVERAL ACTIVITIES THAT WILL ENSURE THAT BOTH THESE REQUIREMENTS CAN BE MET. RESEARCH AND TRAINING IN MANY OF THE DISCIPLINES THAT CONTRIBUTE TO THE DEVELOPMENT AND APPLICATION OF SUPERCOMPUTERS IS FREQUENTLY OVERLOOKED IN UNIVERSITIES DUE IN PART TO LACK OF READY ACCESS. THIS DISCOURAGES GRADUATE STUDENT ENTRIES AS WELL AS RESEARCH IN THE FIELD. LAST YEAR THE UNITED STATES PRODUCED ONLY 225 PH.DS IN COMPUTER SCIENCE AND COMPUTER ENGINEERING. MOST OF THESE DID NOT GO INTO TEACHING, BUT WERE SNAPPED UP BY INDUSTRY. IN THE FIELDS OF COMPUTATIONAL MATHEMATICS, TOTAL PRODUCTION LAST YEAR WAS LESS THAN 30 PH.DS. SINCE THE TRAINING OF SOPHISTICATED SCIENTISTS IS A PROBLEM WITH A CLOSE TO TEN YEAR TIME CONSTANT, NO OVERNIGHT SOLUTIONS ARE POSSIBLE. HOWEVER, NSF HAS BEGUN INITIATIVES THAT WILL ALLEVIATE THIS PROBLEM, AND WE EXPECT DOE AND DOD TO PLAY AN IMPORTANT PART TOO.

JUST AS WE'RE ENCOURAGING GREATER COOPERATION BETWEEN INDUSTRY AND THE FEDERAL LABS IN DEVELOPING FUTURE SUPERCOMPUTERS, WE'RE ALSO ENCOURAGING GREATER LINKS BETWEEN THE FEDERAL LABS AND UNIVERSITIES ON SUPERCOMPUTER RESEARCH. AT THE POINT, FOR EXAMPLE, WHERE NOVEL ARCHITECTURAL IDEAS MUST BE TURNED INTO EXPERIMENTAL MACHINES, THERE IS CLEARLY A ROLE FOR GREATER UNIVERSITY-LABORATORY INTERACTION. ONE EXAMPLE IS THE RECENT INSTALLATION OF A DENELCOR HEP MACHINE AT LOS ALAMOS, "OUTSIDE THE FENCE," AND ACCESSIBLE BY MANY UNIVERSITY RESEARCHERS WITH DOE GRANTS. THIS MACHINE IS A PARALLEL PROCESSOR, AND IT'S

GIVING BOTH LABORATORY AND UNIVERSITY RESEARCHERS THE OPPORTUNITY TO EXPLORE HOW TO PROGRAM FOR A MACHINE THAT HAS BEEN DESCRIBED AS A FORERUNNER OF THE SUPERCOMPUTERS TO COME. ANOTHER INSTITUTIONAL MODEL IS THE RESEARCH INSTITUTE FOR ADVANCE COMPUTATIONAL SYSTEMS AT NASA'S AMES LABORATORY IN CALIFORNIA. THIS JOINT INDUSTRY, UNIVERSITY, AND GOVERNMENT LABORATORY RESEARCH INSTITUTE, AFFILIATED WITH AMES, IS PROVING VERY SUCCESSFUL AT ATTRACTING OUTSTANDING VISITING RESEARCHERS.

AT THE SAME TIME SEVERAL UNIVERSITIES HAVE ACKNOWLEDGED THE NEED FOR GREATER ACADEMIC RECOGNITION OF SCIENTIFIC COMPUTING, NUMERICAL ANALYSIS AND COMPUTATIONAL MATHEMATICS AS A LEGITIMATE SCIENTIFIC FIELD. DOE AND NSF ARE ENCOURAGING THIS TREND THROUGH THEIR RESEARCH SUPPORT IN THIS AREA.

ACCESS TO SUPERCOMPUTERS FOR U.S. SCIENTIFIC AND ENGINEERING RESEARCHERS, WHILE IN MANY WAYS A SEPARATE ISSUE, IN OTHER WAYS IS CENTRAL TO THE DEVELOPMENT OF FUTURE SUPERCOMPUTING SYSTEMS. ACCESS REPRESENTS THE DEMAND SIDE OF THE EQUATION, AND IT HAS GROWN AS THE MACHINES THEMSELVES HAVE DEVELOPED. MANY AREAS OF RESEARCH IN BOTH THE SCIENCES AND ENGINEERING ARE BECOMING TRACTABLE FOR THE FIRST TIME BECAUSE OF THE CAPABILITIES OF CURRENT AND PROPOSED SUPERCOMPUTERS. TODAY, HOWEVER, MOST OF OUR GRADUATE STUDENTS AND FACULTY HAVE NO WAY TO ACCESS THE POWER OF THESE MACHINES FOR THEIR RESEARCH.

THIS ADMINISTRATION EXPECTS TO PROPOSE SEVERAL INITIATIVES WHICH WILL, OVER THE COURSE OF THE NEXT SEVERAL YEARS, ENSURE THAT U.S. RESEARCHERS WITH MERITORIOUS RESEARCH PROPOSALS WILL HAVE ACCESS TO THE NECESSARY SUPERCOMPUTERS AND THAT PH.D'S WILL NOT BE GRADUATED FROM OUR LEADING UNIVERSITIES WITHOUT BEING FAMILIAR WITH THEIR CAPABILITIES. THIS WILL NOT REQUIRE A SUPERCOMPUTER ON EVERY CAMPUS, BUT CAN BE DONE WITH A SMALL NUMBER OF SUPERCOMPUTER CENTERS. THE EMPHASIS THERE WILL BE ON SERVICE TO VISITING AND REMOTE USERS, AND THIS REQUIRES A CRITICAL MASS OF EXPERIENCED OPERATORS AND IN-HOUSE RESEARCHERS. BECAUSE OF THE LIMITED EXPERIENCE IN MOST UNIVERSITIES IN USING SUPERCOMPUTERS AND IN PROVIDING LARGE SCALE USER SERVICES OVER A NETWORK, THESE CENTERS WILL BE ESTABLISHED SLOWLY. THEY MAY GROW, FOR EXAMPLE, FROM AN ESTABLISHED RESEARCH PROGRAM. THEY WILL BE MULTIPURPOSE FACILITIES, PROVIDING SERVICES TO RESEARCHERS, TRAINING GRADUATE STUDENTS, AND STIMULATING COMPUTER SCIENCE RESEARCH ON FUTURE SUPERCOMPUTERS. THESE CENTERS WILL BE LINKED BY SEVERAL, EVENTUALLY COMPATIBLE, HIGH BANDWIDTH NETWORKS. SUPERCOMPUTER ACCESS IS ALREADY BEING ADDRESSED IN FY84 IN BOTH DOE AND NSF, AND WE EXPECT IT TO EXPAND SYSTEMATICALLY AS THE DEMAND EVOLVES AND USERS' EXPERIENCE AND CAPABILITIES GROW TO UTILIZE IT EFFECTIVELY.

THERE HAVE BEEN RECENT CALLS FOR A "CRASH" GOVERNMENT SUPERCOMPUTER

PROGRAM TO RESPOND TO THE JAPANESE CHALLENGE. THERE HAVE EVEN BEEN PROPOSALS FOR A "NATIONAL SUPERCOMPUTER CENTER." WE SHOULD REMEMBER, HOWEVER, THAT EVEN THE JAPANESE HAVE NOT CENTRALIZED THEIR SUPERCOMPUTER DEVELOPMENT. U.S. INDUSTRY HAS ALREADY RESPONDED TO THE JAPANESE RESEARCH CHALLENGE. U.S. PRIVATE FUNDING FOR SUPERCOMPUTER R&D IS UP SIGNIFICANTLY. THE NATIONAL INNOVATION AND PRODUCTIVITY ACT OF 1983, RECENTLY PROPOSED BY THE ADMINISTRATION, WILL SUPPORT AND ENCOURAGE COLLABORATIVE INDUSTRY EFFORTS LIKE THE MICROELECTRONICS AND COMPUTER TECHNOLOGY CORPORATION AND THE SEMICONDUCTOR RESEARCH CORPORATION. INDUSTRY AND GOVERNMENT TOGETHER HAVE ESTABLISHED THE CENTER FOR INTEGRATED SYSTEMS AT STANFORD UNIVERSITY, WHICH IS BEING JOINTLY FUNDED BY BOTH DOD AND ALMOST TWENTY COMPANIES. OTHER COOPERATIVE ACTIVITIES ARE UNDERWAY AT CORNELL, CARNEGIE MELLON, MIT, RENSSELAER POLYTECHNIC INSTITUTE, AND SEVERAL OTHER UNIVERSITIES. THE STRENGTHS OF THE U.S. RESEARCH SYSTEM, WHICH IS UNMATCHED BY ANY OF OUR COMPETITORS, ARE ITS DIVERSITY AND COMPETITIVENESS. THE SUPERCOMPUTER PROGRAMS BEING PROPOSED BY THIS ADMINISTRATION SEEK TO NURTURE AND CAPITALIZE ON THOSE STRENGTHS, ON THE ENTREPRENEURIAL SKILLS OF U.S. INDUSTRY, AND ON THE GOVERNMENT'S OWN CONTINUING REQUIREMENTS FOR SUPERCOMPUTERS.

IN SUMMARY, OUR PROGRAM WILL ENSURE A CONTINUED SUPPLY OF SUPERCOMPUTERS SECOND TO NONE MANUFACTURED IN THE U.S. IN THE SHORT TERM WE'LL DO THIS BY ENCOURAGING THE GOVERNMENT'S CONTINUED ACQUISITION OF SUPERCOMPUTERS, AND BY RE-ESTABLISHING

CLOSE COOPERATION BETWEEN VENDOR AND USER. IN THE LONGER RUN, WE INTEND TO ENHANCE GOVERNMENT SUPPORT FOR SUPERCOMPUTER RELATED RESEARCH IN UNIVERSITIES, FOR GRADUATE TRAINING, AND FOR GUARANTEED PURCHASES OF ADVANCED MACHINES. WE WILL ALSO ENSURE READY ACCESS TO SUPERCOMPUTERS FOR U.S. RESEARCHERS ACROSS THE NATION. THIS COUNTRY HAS ALWAYS THRIVED ON CHALLENGES, AND THE U.S. WILL CONTINUE TO SET THE STANDARD FOR SUPERCOMPUTERS THAT OUR COMPETITORS MUST SEEK TO EMULATE.

MR. CHAIRMAN, THIS CONCLUDES MY FORMAL STATEMENT. I WOULD BE PLEASED TO RESPOND TO ANY QUESTIONS.

Followup Questions and Answers, Association of American Universities and National Association of State Universities and Land-Grant Colleges

SUPERCOMPUTER HEARING QUESTIONS

for AAU and NASULGC

1. From our records it appears that Purdue has become the third American university to obtain a supercomputer, following Colorado State and Minnesota.

   a. Who operates the system? How were they trained?
   b. Who do you serve? Is this a local (campus) network or do you provide the capability to a wider group?
   c. Does the computer center provide programming support to users? If so, how were the support personnel trained? If not, who trains the users?
   d. Please comment on vendor support both for hardware and software maintenance and development.
   e. Please discuss computer center charges and the University's revenue/cost philosophy.
   f. Possibly it is too early in Purdue's experience to answer such a question. However, your comments on computer center efficiency, both now and planned, would be helpful. What unforeseen problems have arisen? What unanticipated benefits have been obtained?
   g. What guidance would you give to other colleges and universities which hope for supercomputer possession?
   h. How many hours per day does the supercomputer operate? What is the present cost per hour for a user?

2. How did Purdue University finance and secure installation of a supercomputer? (They did it without Federal funds). Was it 100% funded by the State of Indiana?

3. You comment in your testimony on the importance of NSF in providing greater access to supercomputers for universities. Yet Purdue University, Colorado State University, and University of Minnesota are three universities that bought supercomputers without Federal help. These actions by universities have placed some pressure on the university system to make these programs work. What then is the proper role for the Federal government (NSF, DOE, etc.)?

4. Is remote access to supercomputer centers adequate for university researchers' needs? What are the impacts of having or not having a supercomputer in an academic institution?

5. What is the federal role in stimulating supercomputer research at universities in hardware architecture design and software application development?

6. Are there plans to develop university centers or consortia for supercomputer use and access? If so, what institutions could be serviced, what are the needs for such a center, and should the Federal Government give consideration to stimulating universities to develop supercomputer consortia?

RECEIVED
11 - 1984
COMMITTEE ON SCIENCE
AND TECHNOLOGY

PURDUE UNIVERSITY

PRESIDENT

December 28, 1983

The Honorable Don Fuqua, Chairman
Committee on Science and Technology

The Honorable Larry Winn, Ranking Minority
U. S. House of Representatives
Rayburn House Office Building, Suite 2321
Washington, D.C. 20515

Dear Congressmen,

Dr. Robert L. Clodius, President of NASULGC, asked me to respond to the two "Supercomputer Hearing Questions" that pertain specifically to Purdue.

Questions:

1. "From our records it appears that Purdue...

   a. The Purdue University Computing Center operates the Cyber 205 as well as the other computing equipment managed by the Computing Center (including a CDC 6500, two CDC 6600 and several Dec Vax 11/780 computers). Soon after we ordered the Cyber 205 in November 1981, we began a training program for the Computing Center staff. Our contract with Control Data Corporation provided they present a series of intensive courses at Purdue on Cyber 205 systems and applications programming.

   b. The Purdue University Computing Center serves the entire University. We also provide service to outside users and have just completed a connection to the Indiana University Computing Center.

   c. By the time the Cyber 205 was installed in January 1983, our user services staff was ready to provide assistance to others in the use of the Cyber 205. A number of major programming projects in the user services department including the development of a SPSS - Statistical System for the Cyber 205 and vectorization of selected IMSL routines, helped our staff develop its own expertise in this area.

WEST LAFAYETTE, INDIANA 47907 317-494-9708

d. CDC support for hardware maintenance has been very good. CDC support for software has been adequate.

e. The current use rate for the Cyber 205 is $750 per machine hour. A machine hour is defined at 3000 billing units. The number of system billing units by a job is equal to the number of central processor (cpu) seconds plus a weighted sum of four input-output parameters, for disk access, disk sectors transferred, small page faults and large page faults. In absence of any input-output activity a system billing unit is a cpu/second. The $750 per machine hour is equivalent to .21 per machine second. Usually the charge for cpu time is 70-90% of the actual charge. The total amount of billable machine time is estimated as 4800 machine hours. Costs are computed based on OMB Circular A-21 procedures. In general we try to obtain 20% of the Purdue University Computing Center costs from sponsored research. (This means about 1/5 of research computing is charged externally, 4/5 internally.)

f. Addition of the Cyber 205 has proceeded smoothly and on schedule. We have no unforeseen problems. We have not received any benefits that were not anticipated.

g. We look at the Cyber 205 as the highest level in our computer network. It is part of the total system and as we become more proficient with interfacing it we will have a total computing system, where more difficult and/or time-consuming computations rise to the highest level. The Cyber 205 is available for classwork, research, undergraduates, graduates, faculty and staff.

h. The production schedule now in effect for Cyber 205 service to users is from 9:00 each morning to 5:15 the next morning. Production starts at 10:00 a.m. on Monday and Wednesday, and at 12:00 noon on Sunday. Production ends at 11:00 p.m. on Friday and at 6:00 p.m. on Saturday. The morning breaks are used for maintenance and for housekeeping (such as disk file backup). The costs per hour were discussed under c. University users costs are paid for by a general computing fund. (Again we attempt to collect 20% of the cost from funded research.)

2. For all practical purposes the Cyber 205 was financed by the State of Indiana through bonding authority. Some internal funds were used to modify space and provide auxiliary equipment.

Thank you for the opportunity to enhance my testimony by responding to your questions.

Sincerely yours,

Steven C. Beering
President

SCB:ms
cc: R. L. Clodius

National Association
of State Universities
and Land-Grant Colleges

Robert L. Clodius, President

January 3, 1984

Congressman Don Fuqua
Congressman Larry Winn, Jr.
House Committee on Science and
Technology
2321 Rayburn HOB
Washington, DC 20515

Dear Congressmen,

I appreciate the opportunity to amplify the testimony that was presented by President Beering of Purdue at the Committee's hearings on SUPERCOMPUTER USE IN SCIENTIFIC AND TECHNOLOGICAL RESEARCH in November. I have asked Dr. Beering to respond directly to you on the questions that pertain specifically to Purdue and in this letter will answer the questions that involve broader national issues of university research and development of supercomputers.

QUESTIONS:

3. "WHAT IS THE PROPER ROLE . . ."

While there are three universities that have made the substantial investments necessary to install supercomputers, and while it is true that those universities did not have direct aid in those investments, in each case it was expected that the Federal government, through research grants and contracts, would make a significant contribution in the funding for the operation of the centers. Similarly, a university will make investments in research instruments in a field, chemistry for example, where usage by investigators with federal grants or contracts is expected to support the operation of the facility.

Each of the universities evaluated the risks, the benefits, and the required resources and boldly went ahead. Each of the universities expected that the addition of a research instrument -- a supercomputer -- would both aid the advancement of science and provide an attractive resource which would draw research activities to its campus.

Some questions that need to be addressed which deal both with the university system and the Federal government are:

1) How many universities can generate and allocate sufficient capital to acquire the supercomputers needed for scientific research use? How many should?

2) Does the system of recovering costs from computers provide sufficient operating funds for the universities to staff and operate supercomputer facilities effectively?

3) What are the Federal strategic options?

The first question, which deals with the capital funds required, can be restated: How many universities can find $20 million to invest in access to a specialized instrument for research purposes. While the Bardon report and others predicted that more than ten to twenty supercomputers would be needed to support university research access in the next five years, those needs are dwarfed by the campus needs for general access to computing and information services. While current expenditures for higher education computing are on the order of two billion dollars per year, it has been estimated that the capital needs to widen the access to information and communication services on each campus will be forty billion dollars over the next five years. This means that the question of allocation of scarce resources for computing will be a crucial issue for universities. Faced with choices between widening demands in all fields and the need for large investments to benefit a relatively small number of researchers, few universities will be able to make supercomputer investments without federal support.

The second question, cost recovery, deals with the impact of federal policies on fairly apportioning costs of use of shared facilities. Without recounting efforts by the universities and the Federal government to agree on revisions to the most recent costs instrument -- OMB Circular A-21 -- I would like to note one aspect of the regulations, that time cannot be given away. This means that student use must be accounted and paid for by the institution.

Supercomputers are shared instruments which are much more than the hardware for the supercomputer. The resources required to support a supercomputer involve:

1) The supercomputer, its fast and slower speed memory (ranging from semiconductor to disk and optical, from tape units to automated tape archives);

2) Communications support facilities to handle users within the installation and remote;

3) Software and necessary systems and user support staff;

4) Operations and maintenance support staff.

It should be noted that DOD and DOE supercomputer installations do not follow the same rules as the universities for cost recovery. This has distorted the perception of costs of university computing operations. Costs at most federal installations are set to recover operating, but not capital, costs.

The third question, about federal strategies, should be addressed by identifying the goals for national supercomputer programs. Some considerations are:

1) Necessary research support should not be limited by facilities when the quality of research and its link to our national goals are high;

2) Our national goals for defense and industrial competitiveness require investments by the Federal government when neither private industry nor our other social institutions (such as universities) can make them.

We feel that the critical role of research in support of our national goals of defense and competitiveness means that we must support the access to supercomputers.

What is the best strategy to carry out this program? The strategy proposed in the NSF Bardon report, which calls for using NSF in an appropriate way for the stimulation of research and support of graduate training, and uses Federal laboratories to complement the university resources, seems appropriate. The emphasis needs to be on a program of access rather than one of operation. While not every university needs to operate a supercomputer, all research universities will have faculty members who need access. Access means the ability to have easy use and appropriate support facilities such as user services, documentation, consulting, etc. Any national program must support these broad objectives through a balance of capital funding for support of effective facilities. Shares of use of the services should be in the hands of the investigators and disciplines and carried out through appropriate peer review processes. Sufficient attention needs to be given to national information networks that would allow access.

4. "IS REMOTE ACCESS TO SUPERCOMPUTERS ADEQUATE . . ."

Remote access to supercomputers can be provided in a variety of ways that can meet some needs well and others poorly or not at all, depending on the problem. The recent NSF report on networks (USERNET: A PRELIMINARY ARCHITECTURAL STUDY AND IMPLEMENTATION STRATEGY, JUNE 1983, FRANK KUO ET AL., SRI INTERNATIONAL) points out some of the problems associated with transmission of data. It is important for the remote user to have interactive facilities for the setup and preparation of problems and also some method for a quick look at the results in order to determine whether or not to have final plots, graphics, etc. Remote access can only work if the user support facilities have hot lines, people on duty at all hours, and good documentation. This useful access also rests on the assumption that the user has been able to spend sufficient training time at the installation to build the competence needed to use the new facilities.

The effect of not having a supercomputer installation at an institution are minimal if the support levels are adequate. The impact of having a supercomputer installation at an institution will depend upon whether it is seen as a facility which is operated for the academic community but supported in some planned manner -- or whether it is just a white elephant.

5. "WHAT IS THE FEDERAL ROLE IN STIMULATING SUPERCOMPUTER RESEARCH . . ."

The current role of NSF and other Federal agencies in support of research in supercomputers, or rather, research in new methods of high speed computation, is appropriate. A distinction must be made between research appropriate to universities and the development activities which are more appropriately done through private industry. Development requires large scale projects that are not appropriate to universities.

6. "ARE THERE PLANS TO DEVELOP UNIVERSITY CENTERS OR CONSORTIA . . ."

Consortia are under investigation by a number of universities already supporting supercomputers and for those considering that method of support for the future. The ideas range from establishing the consortia as a way of stimulating faculty access and thus developing higher rates of use, to consortia between several states and in conjunction with Federal laboratories to assemble the support and use needed to fund the installations. Universities should be encouraged to develop consortia as appropriate agents for resource sharing. NSF should take the lead in providing appropriate incentives for such efforts.

I appreciate the opportunity to address these issues. Please let me know if you have any additional questions.

Sincerely,

Bob Clodius
Robert L. Clodius
President

Followups Questions and Answers, National Science Foundation

Question:

How do the Japanese define a generation of computers versus the United States' definition? What generation of computers is the United States in and what were the previous generations?

Answer:

Dividing the development of computers into periods of history or generations has been characterized in the past by focussing on a particular technical feature. The U. S. and Japan agree on the definition of the first four generations:

- The first generation was based on the use of vacuum tube technology;
- The second generation was based on discrete components and transistors;
- The third generation is based on integrated circuits;
- The fourth generation is characterized by computers using very large-scale integrated circuits (VLSI).
- Generations 3 and 4 are in common use in the U. S. and other parts of the world. The fifth generation refers to advanced research and development. Japan's Fifth Generation Computer Project is a multifocussed approach which considers not only a change in device technology but also a change in design philosophy and in fields of application. The fifth generation computer in Japan would be characterized by a change to "knowledge-based" software with a new computer architecture designed for efficient operation of this software. This new system would include an inference and pattern recognition function, a knowledge-base function, and man-machine interfaces that are natural to human communication. In the U. S., the term "Fifth Generation" is not widely used except in reference to the Japanese project. Current U. S. research and development is following several directions including:

  a) The DARPA Strategic Computing Project which also emphasizes the technology and design of high speed "knowledge-based systems";

  b) Multi-processor high-speed vector computers being developed by Cray Research, Inc. and ETA Systems, Inc. with private support;

  c) Other novel high-speed architectures being developed by Denelcor and IBM with private support and by Goodyear Aerospace Corp. with NASA support;

  d) Academic research in highly parallel computer systems, some being conducted with NSF support and some with DOD or DOE support.

Question:

In your view. why are the present university supercomputer centers under-utilized?

Answer:

Several factors combine to lead to the present under-utilization of the university supercomputer centers. First, the high cost of supercomputers, even after substantial educational discounts by the computer manufacturers, stretches the financial resources of most universities to the limit. As a result, university systems are often underconfigured for many problems of scientific interest. Second, the universities must recover the cost of the equipment and operation from the users of the computers. This leads them to charge approximately $2000/hour for supercomputer time, inexpensive by comparison with standard commercial rates but a large expense for researchers who may need hundreds of hours to conduct effective research projects. Individual researchers are reluctant to experiment with this level of cost unless special assistance can be provided. This is particularly true if a major software conversion effort is involved, as is typically the case to efficiently utilize a vector processor. Third, experience has shown that a system of support that depends primarily on payment by researchers for computer costs does not work well. This situation has been aggravated by the large amounts of computer resources that academic scientists and engineers are currently able to obtain without expenditures of grant funds (i.e., the so-called free or university-subsidized computer time). As a result, the financial resources available to support university supercomputer centers are currently insufficient to maintain even a few centers at a level of activity and stability necessary for high quality services.

It is partially for these reasons that the NSF is currently planning multi-year grants or contracts to acquire large blocks of time from computer centers and will allocate the time directly to researchers who need it. This will provide a more stable financial basis for computer centers which will help them plan, equip, and staff their services better.

Question:

What is the most effective approach NSF can take towards making supercomputers available to university researchers? For instance, does NSF plan on supporting supercomputer centers on campuses?

Answer:

For FY 1984 NSF plans to acquire large blocks of time in response to proposals from institutions which are prepared to provide it. NSF will then allocate this time to researchers who request it in proposals. Proposal evaluations and allocations of computer time will be guided by peer review.

As a first step in implementing this plan, the Foundation issued a program solicitation in October, 1983, for proposals to provide computer time and other services. All non-federal organizations were eligible to submit proposals including for-profits, not-for-profits, and academic institutions. Federally Funded Research and Development Centers were also eligible to compete but only with the approval of their parent agency. Proposals were due December 31, 1983, and will be evaluated in the next few months.

Also in October, NSF sent a letter to all grantees inviting them to submit proposals for supercomputer services needed for their research. These proposals, too, will be evaluated in the next few months. The amount of computer time the Foundation will ultimately acquire as a result of these actions will be determined by the need and by the $6.0M budget established by the NSF Director for FY 1984.

Subject to evidence of need and the availability of funds, NSF plans to issue future program solicitations for proposals from institutions which are not now but may wish to become providers of supercomputer services. This solicitation will also be open to organizations of all types, including academic institutions.

The Foundation recognizes the need to maintain flexibility while we determine the best approach for supporting scientific supercomputer centers, based on an understanding of needs and expectations.

Question:

In the best of all possible worlds, should every major research university have a supercomputer on campus? Is remote access to supercomputer centers adequate for university needs? What are the impacts of having or not having a supercomputer in an academic institution? What would the impact of not having access to supercomputers be on the research done in scientific disciplines which are not historically computer related?

Answer:

Supercomputers by definition are the most advanced and fastest computers available at any given time. Thus they are a scarce and expensive resource. During the decade from 1971 - 1982, no university had acquired such state-of-the-art systems, and few are likely to have such research tools of their own in the future. Rather, as with other expensive research equipment, supercomputers will be installed mainly as centralized facilities at a few locations.

Remote access to supercomputers is a realistic option with existing telecommunication technology. For a majority of users, having remote access to a high-quality, service-oriented facility would fully enable them to tackle ambitious new research projects. For some applications requiring transfer of vast amounts of data or high-resolution graphics, high-bandwidth communications links may be needed. On the whole, a substantial fraction of large-scale academic computational research could be carried out given the combination of improved remote access and strengthened support for ancillary local computing equipment.

The choice of location for supercomputer centers involves a variety of issues. There are important reasons to consider placing some of the new facilities in universities and other equally cogent reasons for considering the National Laboratories. The latter have developed considerable experience in meeting the needs of a community of users, as well as providing remote access. University-based supercomputer facilities will certainly benefit from the experience of the National laboratories, and some universities are planning to closely model their operations on the most successful laboratories, using identical hardware and software. The presence of a supercomputer on a university campus can have the principal benefit of exposing students to all aspects of supercomputer research and applications. Universities are catalysts for innovation, and placing the best available tools in such a milieu will lead to many new ideas. These benefits are likely to be lost if all centers were to be off-campus.

Scientists in fields which are not historically computer-dependent are rapidly finding more and more computationally demanding problems, the need for simulation, and the requirements of storing and analyzing large scientific data sets. From mathematics to archaeology, university faculty are finding exciting new applications of this technology. Denying these scientists access to supercomputers may artificially depress progress in these fields, just as research in many other areas has been redirected in the past by the lack of appropriate research tools.

Question:

What would be the major advantages of developing university consortia to operate supercomputer centers? Does NSF have plans to stimulate the development of university consortia for this purpose?

Answer:

Supercomputer centers are costly to establish and operate. Thus only a limited number can be funded. A university consortium has been a successful manager of the NSF-funded computing center at the National Center for Atmospheric Research (NCAR). The Foundation recognizes that this consortium has been instrumental in incorporating the participation of their research community in the allocation of computing resources, setting of plans and priorities, and reviewing of policies and operating procedures. Consortia are being created in the northeastern part of the U.S., in southern California, and other areas around the Nation to operate existing and planned supercomputer centers. The advantages of consortia over other forms of management are that they bring together a broad range of experience, ideas, and management styles, and offer broad opportunities for coordinating research and training activities.

Question:

Several Federal laboratories receive high marks for the quality of the supercomputer services which they provide. What characteristics do you feel have contributed to their success? What would be required for a university supercomputer center to develop the necessary components and environment for quality services?

Answer:

There are numerous aspects that have contributed to the success of the scientific computing centers at several of the national laboratories. These include:

- Stable and adequate funding sources;
- Interactive computing environment on either front-end computers or on the supercomputers themselves;
- Well-configured systems including a hierarchy of storage devices and a variety of output devices;
- Large consultant groups which specialize in scientific and engineering applications and which have experience on a full range of hardware;
- On-line access to large scientific data sets and to the best numerical and utility software libraries;
- Regularly scheduled training programs for new users and advanced training courses for all users;
- Competent management that seeks advice from the user community.

There is one important area the Foundation recognizes that the national laboratories have had little opportunity in which to participate -- the training of graduate students in using and operating supercomputers. For addressing this, we look to the universities. It is not expected that universities, at this moment, can serve the research community with all of the same high-quality services as national laboratories. They have only just begun to acquire supercomputers. However, with competent staffs, development of university consortia, assistance from the national laboratories, adequate funding, and cooperation between industry, government, and universities, we can expect universities to have, in the near future, highly successful advanced scientific computing centers.

Question:

Considering the general computing needs of academic research, what is the best mix of NSF's funds -- upgrading university computing centers; granting dedicated machines to individual researcher's projects; development of national and regional networks; creating supercomputer centers?

Answer:

The appropriate balance of these important ingredients was considered in a study last spring by NSF staff and continues to evolve as the Foundation receives additional input from the research and engineering community. Based on experience in FY 1984, the newly formed Advisory Committee for Advanced Scientific Computing will review the issue of balance at one of its early meetings in order to establish sound policies for FY 1985 and beyond. The best mix of funds may vary from discipline to discipline and can be expected to change as a more complete understanding of the real need is identified; this will require continuing assessment by the Foundation.

Question:

What level of supercomputer research in architecture design and software applications development is NSF supporting?

Answer:

In FY 1984, the Foundation anticipates support for architecture design and software development will increase by about $1.3M to an approximate total of $8.M per year. This will be provided through its Division of Computer Research, the Computer Engineering Program of the Division of Electrical, Computer and Systems Engineering (ECSE), the Division of Information Science and Technology, and the Division of Mathematical Sciences. Most of these funds will be used for software development, but approximately $2.0M will be invested in architecture research.

In FY 83 the Division of Computer Research supported 18 projects totaling about $1.7M and in addition supported four coordinated research activities totaling about $3M (Yale, Rice, Wisconsin and Maryland). These activities included parallel numerical algorithms, large scale numerical computations, and new concepts for hardware development. The ECSE Division supported 24 projects totaling $2M for research related to advanced supercomputers, new languages, and algorithms.

Question:

Prior to 1970, computer research and development benefited from strong interaction between government, universities and industry. Does NSF have plans to re-establish this interaction?

Answer:

Prior to 1972, NSF had an Institutional Computing Services program which helped academic institutions aquire modern computing facilities for research. This program stimulated the use of computers by academic institutions for both research and education and helped spur the development of scientific computers by the computer industry. Today, the Foundation believes that the role of the Federal Government should be different. NSF has identified four areas which require increased support: access to supercomputers, local hardware facilities, enhancement of networking capabilities, and research on new computer architectures and algorithms. To carry out this program NSF will coordinate with other Federal agencies to help ensure appropriate sharing of resources, to avoid duplication of effort, and to maintain a dialogue on policy and technical issues. Also NSF and the computer industry will interact in order to understand industry's perspectives on scientific computing, to be aware of industry's need for university participation in computer research activities, and to coordinate funding and administration of cooperative programs between industry, universities and NSF.

In the area of computer architecture research, however, NSF is actively exploring new collaborative relationships among universities, industry, and government. Very Large Scale Integrated Circuit (VLSI) technology has opened fruitful avenues for academic research in computer architecture. Industry, in turn, not only has interesting research ideas but, also, unique facilities and engineering know-how which are needed by university researchers. As a result, industrial researchers may actively collaborate in the creative aspect of research with university scientists and also provide their unique facilities and experience.

Question:

Also prior to 1970, NSF supported computer centers to improve computing services on campuses. This entire program, from 1956 to 1970, cost about $200M in 1982 dollars. Is there a need for reinstating such a program for supercomputers? What funding level would be necessary to establish a viable supercomputer center program analogous to this earlier effort?

Answer:

There is a need to establish an NSF program to improve access to supercomputers by academic researchers. Steps are being taken in FY 1984 to address this problem. These include the following:

- Initiate a competitive solicitation to request computing time from centers;
- Inform the research community that NSF would welcome proposals requesting access to supercomputers and other large-scale computer systems;
- Determine the need for networks linking researchers with each other and with advanced computing centers;
- Establish an advisory committee to assist the Foundation in evaluating the responses to the solicitation and to advise NSF on policies and plans relative to advanced computation;
- Award funding to several centers for computing time and plan for increasing access in the future.

We must evaluate the need for supercomputer time and facilities in competition with other research needs, e.g., other instrumentation, and general research support. The FY 1984 plans call for the development of a more complete assessment of needs. We will seek advice from the newly established Advisory Committee for Advanced Scientific Computing on funding needs and appropriate NSF program structure.

Question:

Is NSF, in cooperation with other agencies such as DOE and DOD, establishing a mechanism to trade computer time between agencies to permit a researcher who has contracts with more than one agency to use only one computer system, thus saving time and effort?

Answer:

The Foundation is committed to the highest level of cooperation with other agencies in the provision of access by researchers to advanced computing facilities. The recent FCCSET reports on supercomputers, prepared with the full participation of NSF, DOE and DOD personnel, recommended that a mechanism for trading computer time be established. We see no reason not to accept this recommendation. The Foundation is in the process of acquiring large blocks of computer time for allocation to researchers. As soon as the time is acquired, NSF will, together with the other agencies, develop a mechanism that maximizes the ease of access by researchers and minimizes the paperwork. We are also very confident that an appropriate interagency plan will be devised for access not only to numerical supercomputers of the Class VI and Class VII type, but also to experimental computers.

Question:

In FY 84, NSF proposes to make supercomputer time available to their grantees by collecting CPU time at various supercomputer installations and adding that time to research grants. Are any of the three universities with supercomputers contributing time? If not, why?

Answer:

The program solicitation described in response to question 3 above was open to universities as well as other organizations. Indeed, as your question anticipated, proposals were received from universities and commercial organizations with supercomputer facilities. The decision about which proposals to support will be determined with help of peer review taking into consideration such factors as adequacy of facilities, ease of access, availability of supporting services, communications services, and other factors governing the suitability of the supercomputer center to serve scientific research needs.

Question:

In the NSF draft report "A National Computing Enviroment for Academic Research," it was recommended that 10 supercomputer centers be established. Have you set criteria that a host institution should meet to be considered an effective site for a supercomputer and if so, what are those criteria.

Answer:

Yes. In the NSF project solicitation document that requested proposals from centers, the following criteria were listed in decreasing order of importance, with particular emphasis on the first five:

- Quality and reliability of services, including the ease of user access;
- Variety, quality and versatility of systems and applications software packages and the availability of data bases;
- Nature and quality of other research activities at the center, including involvement of students and availability of library facilities;
- Effect of the proposed arrangements on the infrastructure of science and engineering;
- Ability to deliver services in Fiscal Year 1984;
- Cost of computer services;
- Suitability of location of the site with regard to ease of accommodation and transportation for prospective users.

Question:

In the past, supercomputer vendors have provided limited software. The Federal laboratories, especially DOE laboratories, have developed the necessary software and operating systems. If NSF supports supercomputer centers at universities, what will the NSF role be in assisting the centers with obtaining software and developing interactive operating systems to provide high-quality service.

Answer:

The Foundation will support supercomputer centers to provide high-quality computational services to the scientific and engineering community. Hence, not only hardware, but the very best and latest software must be provided to these centers. NSF will provide the funds needed for these purposes. The Foundation is not requesting that these centers participate in the development of systems software; i.e., operating systems, compilers, etc; this is an activity that requires very specialized personnel. Our plan relies on the centers obtaining this software from commercial vendors and other sources such as the national laboratories. Undoubtably, some NSF funded supercomputer centers will conduct research on advanced software, but it is not expected that they will assume the role of either the national laboratories or of commercial vendors in this area.

The situation in applications software, on the other hand, is very different. Universities have contributed greatly to this area throughout the history of computing, including the recent history of fast, vector computers. We expect that contribution to be increased many fold as faculty and graduate students gain access to supercomputers and learn to apply them to their research needs.

Question:

Dr. Gibbons of OTA indicated in the testimony he submitted that to "develop real centers of expertise, the (supercomputer) facilities need to focus on areas of research that share similar problems and research approaches." Could this approach tend to isolate specific disciplines and reduce the benefits of sharing data and techniques over a broad range of disciplines?

Answer:

There are clear advantages to organizing a specific supercomputer facility with an orientation toward a single discipline, and there are equally compelling reasons favoring a cross-disciplinary approach.

A discipline-specific orientation would concentrate scientists with common research interests at a single center. This would encourage scientific interaction and mutual assistance because of the commonality of background and interests. It would encourage the development of shared data bases and specialized equipment. It would also simplify the problems of peer review and oversight. NCAR experience with this mode of operation has been quite good.

On the other hand, a multi-disciplinary facility has its own strengths. Throughout all large-scale scientific computing, there is a common mathematical and computational methodology available to promote significant cross-fertilization between scientists from quite different fields. An inter-disciplinary organization fits well in a university environment where students can play an important role in the use of such facilities. Such considerations have led Cornell University to initiate an inter-disciplinary large-scale computing facility, cutting across all the theoretical sciences.

In a national research effort, both modes of organization can make important contributions. It is, therefore, important to have a variety of facilities and a flexible policy which allows scientists in various fields to choose whichever organizational structure best suits their needs.

Question:

What has NSF done to assess the effectiveness and potential of CSNET? What are NSF's future plans for CSNET?

Answer:

NSF formed several ad hoc review panels during the development of CSNET to assess progress and provide advice about development plans. The network has been in operation now for two years with a steadily growing membership. The computer research community has been enthusiastic about its usefulness. Last year the Foundation instituted a system of membership fees and utility charges for use of the network. One testimonial for the network's effectiveness is the willingness of the users to pay the associated costs. Both academic and industrial users have demonstrated their interest by paying their bills. NSF's direct support of CSNET has been reduced as income from member charges has increased, and it is planned that CSNET will be self-supporting by FY1986.

Until the summer of 1983, CSNET was managed by NSF staff with contracts to several universities and other organizations for software development and network operations. In the summer of 1983, NSF contracted with the University Corporation for Atmospheric Research (UCAR) to manage the continued operation and further development of CSNET with a view toward spinning CSNET off as an independent organization at a later time, if that is deemed desirable by the computer research community. UCAR has formed an Executive Committee of knowledgeable people from the computer research community to advise it. A search for an Executive Director is now being made.

NSF considers the expense of belonging to and using networks to be an appropriate element of cost in grants for support of computer research. As a result, it will support some CSNET costs through research grants in computer research. Also, it plans to support the network costs of some institutions which are primarily undergraduate institutions as a way of helping them integrate their computer science faculty more closely into the mainstream of computer research.

Followups Questions and Answers, Department of Defense

Ada Language

Question: How completely has the software language Ada been assimilated by the respective Services? Please comment on a Service-by-Service basis (Air Force, Navy, Army).

Answer: Each Service is aggressively planning for Ada use. The Army is developing the Ada Language System (ALS). The Air Force is developing the Ada Integrated Environment (AIE). The Navy is building directly on the Army's work and developing the ALS/N (Navy). These efforts are being coordinated through the Ada Joint Program Office so as to meet immediate needs for the technology in each Service, to transfer knowledge and technology, and to promote commonality. The Services have identified thirty (30) system programs which will use Ada in the near term (next three years). Ten (10) of these are Army programs, eight (8) are Navy programs, and twelve (12) are Air Force programs. After this transition period, Ada will be the norm and will be expected to be used on all systems.

Symbolic Processing and Vector Processing Funding

Question: What percentage of total funding in DoD is directed to symbolic processing and to vector processing? Do you anticipate that both processing philosophies can be incorporated into one system?

Answe Symbolic Processing deals with non-numeric information and the relationships between pieces of non-numeric information, whereas vector processing deals with computations involving lists or arrays of numeric values. The two types of processing essentially address different kinds of problems. DARPA's Strategic Computing Program primarily focuses on symbolic processing. It is budgeted at $50 million (M) in fiscal year (FY) 1984, $95M in FY 1985, and $150M in FY 1986, FY 1987, and FY 1988. The only vector processing development effort in DoD is the S-1 Program jointly funded with DOE. It is currently budgeted at approximately $3.5M annually from each Department from FY 1984 through FY 1988. DoD's needs for vector processing machines have essentially been met by industry.

It is believed that symbolic processing and vector processing are not likely to coexist in the same system. Efficient information addressing mechanisms for symbolic processing appear to be inconsistent with data addressing for vector processing. However, additional research and development is required to determine the extent of such incompatibilities and what compromises might be appropriate. The DARPA program will investigate this issue and possible resolutions.

Computer Classifications

Question: What is the class 6 computer and kindly describe the characteristics of other classes of computers. Could you provide the Committee with the Japanese classification of computers versus that of the United States?

Answer: The term "class 6 computer" has a very loose definition created by the Department of Energy (DOE), but it has been widely adopted to represent machines in the performance range of the Cray-1 and the Control Data Cyber 205. DOE first coined the term "class 4 computer" to establish a generic name for computers that have performance levels on the order of the Control Data 7600 and the IBM 360/195. As DOE acquired and measured other machines with approximately the same capability, they too were designated class 4 machines. Later as new systems with higher performance levels became available the DOE created class 6 for machines with approximately four times the performance of class 4 machines. DOE does not designate a classification until after a machine has been installed, DOE's programs have been optimized to take advantage of its unique characteristics, and its actual performance measured. Class 5 was next established and includes the IBM 3033, although other machines exist that would probably fall within the same performance range. Class 7 will be used to represent machines that have six to ten times the performance of class 6 machines. It is expected that the Cray-2 will be designated a class 7 machine.

The Japanese designation "Fifth Generation" is a chronological classification and is meant to infer the next major new generation of computers. The first generation of commercial computers employed vacuum tubes. The second generation employed transistors. The third generation employed integrated circuits and the fourth generation employs Very Large Scale Integration or VLSI technology. Neither the U.S. nor the Japanese have established a well defined performance classification scheme for computers.

The Role of Large-Scale Computing in National Security

Question: What part does large-scale computing play in national security and what is DoD's responsibility in fulfilling this role.

Answer: Large-scale computers are indispensible in a number of areas critical to our national security. They are necessary for nuclear weapons analysis and design, weather forecasting, aerodynamics and aerospace research, large simulations and modeling, communications security, intelligence processing, and for supporting scientific computation in a large number of activities associated with defense systems development. They give U.S. scientists an additional edge and thus they play a very important role in our efforts to maintain military leadership. DoD is responsible for obtaining the most advanced computers available to meet these needs. The U.S. computer industry has made tremendous advances over the past two decades and has been able to meet DoD's large-scale computing needs.

Number-Cruncher vs. Artificial Intelligence Research Needs

Question: How do your supercomputer needs compare with DARPA's? What DoD mechanisms exist for coordinating research needs between classic number-cruncher supercomputers and artificial intelligence machines?

Answer: DARPA's Strategic Computing program is essentially building the technology base to develop intelligent machines for advanced military applications. The Science and Technology Program managed in the Office of the Deputy Under Secretary of Defense for Research and Advanced Technology does not include a counterpart program for advanced number-cruncher supercomputers. The technology base for developing such machines to a large extent already exists in the U.S. computer industry. Advancing the state-of-the-art does not necessarily require direct government funding. Other approaches have been proposed which could provide industry with the revenue to support the necessary research. These include additional government purchases of the current generation of such machines to reduce the economic risk for the industry. Coordination of research needs between the two types of requirements is performed through both formal and informal reviews involving representatives of all the appropriate organizations.

Level Of Supercomputer Use

Question: What is the present level of supercomputer use in the DoD? What is the projected future use?

Answer: The Air Force currently has a Cray-1 at the Air Force Weapons Laboratory and has procured a Cray-1 for integration with the wind tunnel facility at the Arnold Engineering Development Center. The Navy has a Control Data Cyber 205 at the Naval Fleet Numeric Oceanographic Center, and has initiated procurement for a supercomputer at the Naval Research Laboratory to replace their Texas Instruments Advanced Scientific Computer. The Number and type of supercomputers at the National Security Agency is classified. In addition the Naval Ships Research and Development Center, the Army's Ballistic Missile Defense Program, the Air Force Global Weather Central, and the Air Force Aeronautical Systems Divison have plans for such computers in the near future. The Arnold Engineering Development Center also plans to procure an additional supercomputer in the near term.

Utilization of Existing Supercomputers

Question: It has been suggested that the supercomputers in use by the DoD are underutilized? If this is the case can you explain why?

Answer: This is not the case. In general the supercomputers within DoD are heavily loaded. Procurement of additional machines is planned since existing facilities cannot handle the growing load. One machine has been delivered to an Air Force contractor for integration with a wind tunnel facility and has not been delivered to the government. This machine is not yet being used. Suggestions have been made that this machine could be used for other applications while awaiting completion of the wind tunnel facility. However, the configuration of this machine does not lend itself to the general scientific computation problems suggested. The memory size is inadequate for the normal support software required for this use and temporary modifications to make such provisions are not justified. Another machine is being procured for the same laboratory to meet the non wind tunnel requirements.

Supercomputer Access for Researchers

Question: What arrangements are in place for access to DoD supercomputers for researchers in academia and the private sector? What future arrangements are planned?

Answer: DoD is very much concerned about access to supercomputers by research contractors, particularly in academia. DoD and other agencies have been addressing some of these concerns under the auspices of the Office of Science and Technology Policy. Specifically, the problem of providing universities with access to supercomputers, providing training opportunities for students, developing algorithms for supercomputers, and determining the proper role of government in the development of supercomputers for the future are the issues being addressed.

Currently the U.S. Army now has arrangements for university researchers to have access through their Seromechanics Laboratory to the Cray 1 supercomputer at NASA's Ames Research Center, as well as providing access to the Control Data 7600 at their Ballistics Research Laboratory. The Navy is budgeting one million dollars per year for computer time on the planned new supercomputer for the Naval Research Laboratory. The Air Force Office of Scientific Research (AFOSR) purchases time for universities on the Cray supercomputer at the Air Force Weapons Laboratory. AFOSR also provides universities with grants for computer services at commercial facilities and government sponsored laboratories such as Los Alamos.

ARPANET And CSNET

Question: DoD manages a network, the ARPANET, used by government, private and academic researchers working on defense projects. Does DoD plan to expand the use of that network to provide access to a larger audience of researchers? Do you see NSF's network, CSNET, as a duplication of the services provided by the ARPANET? Do cooperative arrangements exist between CSNET and ARPANET to provide researchers, serviced by each network, with greater access?

Answer: The ARPANET is an operational DOD network and is not intended to compete with comparable commercial services. Accordingly, before ARPANET service is provided to any non-U.S. Government activity, it must be determined that adequate commercial service is not available. The ARPANET is intended to be used solely for the conduct of or in support of official U.S. Government business. DOD does not plan to provide access beyond this audience.

Recently DoD logically divided the ARPANET into two separate networks. One has been named MILNET and is primarily for operational military use, the other retains the name ARPANET and is used as a research network. A non-Government activity must have a U.S. Government activity acting as a sponsor. An application for an ARPANET subscription must clearly state how the proposed service is in the best interest of the U.S. Government, is essential to mission fulfillment, does not violate privacy laws, and why adequate commercial service is not available. The contract number of the contract between the sponsoring activity and the non-Government activity requiring access to the ARPANET may be requested.

CSNET was developed to provide services to communiations researchers and industrial users not funded by DoD. It is intended to be self supporting, i.e., not dependent on federal subsidies, it is a logical network made up of a number of public physical networks and DoD's ARPANET. Communication gateways exist between the public at works and ARPANET, but general access to ARPANET computers is not provided. Physical capability exist to permit such access on a very limited case-by-case basis a few nonARPANET users have been given access to ARPANET computer through the gateways. Indeed CSNET is not duplicative of ARPANET services since it provides general digital network facilities to nonDoD researchers and users. DoD does not plan to extend the use of the ARPANET to provide access to a larger audience of researchers.

Consolidating Computer Time Requirements

Question: Is DoD, in cooperation with other agencies such as DOE and NSF, establishing a mechanism to trade computer time between agencies to permit a researcher who has contracts with more than one agency to use only one computer system, thus saving time and effort?

Answer: No formal mechanism has yet been established however DoD will look into making such provisions with other appropriate agencies.

Additional Material Submitted for the Record

REPORT
OF THE
PANEL ON LARGE SCALE COMPUTING
IN SCIENCE AND ENGINEERING

Peter D. Lax, Chairman

Under the sponsorship of

Department of Defense (DOD)
National Science Foundation (NSF)

In cooperation with

Department of Energy (DOE)
National Aeronautics and Space Administration (NASA)

December 26, 1982

This report does not necessarily represent either the policies or the views of the sponsoring or cooperating agencies.

I. Introduction

Supercomputers are the fastest and most powerful scientific computing systems[1] available at any given time: they offer speed and capacity, or special characteristics, significantly greater than on the most widely available machines built primarily for commercial use.[2] Large scale scientific computing is the application of supercomputers to the solution of a model or simulation of a scientific or engineering problem through the appropriate use of numerical algorithms and techniques.

The availability of supercomputers during the past thirty years has been crucial to the Nation's advances in science, engineering, national security, and industrial productivity. Supercomputers have been essential to scientific and engineering investigations in areas such as atmospheric research, astrophysics, molecular biology, integrated circuit design, and fusion research. The weapons programs of DOE, cryptographic analysis, and weather forecasting are dependent on the availability of computational facilities. The use of supercomputers in the aerospace, petroleum, semiconductor, and nuclear industries contributes substantially to the nation's productivity. The development of supercomputers has significant spinoffs for all the technologically based components of the national economy. Research and development in semiconductor technology and in computer research has directly supported and expanded the defense, industrial, medical, and consumer segments of the economy.

The U.S. is the acknowledged leader in the development and use of supercomputers. In 1970 this Nation was preeminent in all aspects of electronic, computer, and computational technology. However, America's present leadership in supercomputers is challenged in the areas of components, development, and applications. Recently, Hitachi has begun marketing what is claimed to be the first 16k bipolar ECL RAM[3]; this device, representative of the continuing advances of Japanese microelectronic manufacturers, is designed for applications in today's scientific computers. Fujitsu, Nippon Electric Company, and Hitachi have each developed supercomputers, which are claimed to compare favorably with the available, or announced, American systems. American universities and research centers, which have historically created new applications of supercomputers and sustained research in computational mathematics, have lagged behind their Japanese and European counterparts in the installation of supercomputers.

---

[1] These include specialized machines, such as array processors, that are equal to or, for some problems, more powerful than general purpose mainframes.

[2] The term supercomputer, as used in this report, encompasses hardware, software, supporting peripherals, and the facilities and personnel needed for their appropriate use. Appendix I reproduces three papers presented to the Panel: NCAR Computing Capabilities and Services, by W. Macintyre; Magnetic Fusion Energy and Computers, by J. Killeen; and The Potential of Los Alamos National Laboratory to Provide Large Scale Computational Capabilities to the Research Community, by B. Buzbee and D. Sparks. These papers present descriptions of supercomputer facilities and of their access from remote facilities through networking.

[3] See Electronic Engineering Times, December 6, 1982.

Significant national thrusts in supercomputing are being pursued by the governments of Japan, West Germany, France, and Great Britain. Some of these, notably the Japanese effort, center on the development of supercomputers; others, on the provision of supercomputers, or access to them through networks, to the scientific and engineering research community. The British program, for example, is designed to provide research scientists and engineers in academic and government laboratories access to supercomputers through modern workstations connected to a high-speed national network. These agressive foreign national initiatives provide a striking contrast to the current state of planning in the U.S.. The domestic computer industry continues its vigorous research and development efforts in the supercomputer field; however, it is felt that these efforts, necessarily dictated by commercial conditions, are less than they could be and far less than should be for the national scientific and technical capability as a whole. The U.S. research community does not have sufficient access to supercomputing facilities, as is documented in numerous studies, papers, and reports directed toward specific disciplines and specific agencies. A partial bibliography of these studies is included in this report.

Expressions of concern that the U.S. is failing to exploit its position of leadership in supercomputing are being voiced from many quarters. Reflecting this concern, the NSF/DOD Coordinating Committee requested, in April of 1982, that a workshop be organized to explore the problems, needs, and opportunities in large scale computing. This Workshop, sponsored by NSF and DOD with the cooperation of DOE and NASA, was led by a panel of fifteen scientists and engineers from a broad spectrum of disciplines. It took place at the NSF on June 21-22, 1982, and was attended by over one hundred participants. Experts in the use, design, and management of large scale computers from the computing, defense, and other industries, government laboratories, universities, and research funding agencies were included. The lists of the participants in this Workshop are contained in the Supplement.

The Panel assessed the role of supercomputing in scientific and engineering research; surveyed the current use, availability, and adequacy of supercomputers; and considered near- and long-term needs. Subsequent to the June 21-22 Workshop, numerous meetings of smaller groups of participants have taken place; in particular, experts on computer development (Group 3 of the Lax Panel) met at Bellaire, Michigan, on August 23-24, 1982, to further explore avenues for assuring the development of future supercomputers. From these meetings a large number of suggestions and position papers have been directed to the Panel and to the Organizing Committee. This report is an attempt, on the part of the Organizing Committee and the Panel, to outline both the results of the Workshop and the subsequent discussions and contributions. The Panel has chosen not to repeat all the detailed technical arguments or examples of the use of supercomputers found in the literature. A bibliography and appendices are included.

ssues involve many Federal agencies, government laboratories, niversities, private sector companies, and scientific disciplines, they eed to be addressed on a National basis and require Federal study, lanning, and support. The Panel's report attempts to bring the undamental issues to the attention of policymakers; however, it eliberately avoids details of an organizational, programmatic, or udgetary nature.

I Summary of Findings and Recommendations

ummary of Findings

arge scale computing is a vital component of science, engineering, and echnology[4], bringing together theory and applications. It is essential or the design of many technologically sophisticated products, and is aking possible for the first time the analysis of very complex scientific nd engineering problems which to date have defied analytical and xperimental techniques. Examples of the importance of supercomputing are riefly noted below[5].

Renormalization group techniques[6] are a major theoretical breakthrough that provide a new framework for the understanding of a number of unsolved scientific and engineering problems ranging from problems in quantum field theory, the onset of phase transitions in materials, the development of turbulence, propagation of cracks in metals, and the exploitation of oil reservoirs. Only a minute fraction of these problems can be solved analytically. Large scale computational techniques have been essential to the use of renormalization group methods, and even today's largest computational machines are not sufficiently powerful to address most of these problems[7].

Aerodynamic design using a supercomputer has resulted in the design of an airfoil with 40% less drag than that developed by previous experimental techniques[8]. The solution of the Navier-Stokes equations with sufficient

---

See, for example, R & D for National Strength, Center for Strategic and International Studies, Georgetown University, Washington, D.C., 1982. See also The Defense Science Board Summer Study, 1981.

In the Supplement of this Report, H. B. Keller and J. R. Rice describe in some detail scientific and engineering areas in need of supercomputers. Appendix II contains a number of examples of scientific and engineering problems successfully addressed on supercomputers as well as additional examples requiring supercomputing capabilities not yet available.

These techniques were devised to handle the description of phenomena with interactions spanning an extremely wide scale. K. G. Wilson was awarded the 1982 Nobel Prize in Physics for his contributions to the theoretical development and application of renormalization group techniques to critical phenomena. See his contributions to this report in the Supplement, as well as in Appendices II and III.

See, in Appendix II, several problems posed by K. G. Wilson.

For a detailed description see "Trends and Pacing Items in Computational Aerodynamics", [40].

resolution to represent faithfully fluid behavior became practical with the current (Class VI) generation of supercomputers. The wings of the Boeing 767 and of the European Airbus 310 were designed by computational methods on such machines, resulting in this most significant improvement.

The aerodynamic design of an entire aircraft is not feasible with today's supercomputers; it is estimated that machines 100 times more powerful are needed for this purpose. The design of jet engines, involving the simulation of complex three-dimensional fluid flows and associated chemical reactions, also requires significantly increased computational capability and capacity.

In design, especially of advanced weapons systems, large scale computational modeling is an essential substitute for experimentation. Similarly, the design of future generations of nuclear power plants, and their operation--relying on real-time simulation for their control--require computational facilities several orders of magnitude greater than those available today.

Perhaps the most significant applications of scientific computing lie not in the solution of old problems but in the discovery of new phenomena through numerical experimentation; the discovery of nonergotic behavior, such as the formation of solitons, and the presence of strange attractors as universal features common to a large class of nonlinear systems are examples of this scientific process.

Current and feasible supercomputers are extremely powerful scientific and engineering tools. They permit the solution of previously intractable problems, and motivate scientists and engineers to explore and formulate new areas of investigation. They will surely find significant applications not yet imagined. For these reasons, the Panel believes that it is in the National interest that access to constantly updated supercomputing facilities be provided to scientific and engineering researchers, and that a large and imaginative user community be trained in their uses and capabilities.

The U.S. has been and continues to be the leader in supercomputer technology and in the use of supercomputers in science and engineering. The present position of leadership is evidenced by the dominance of the supercomputer market by American producers and by the successful exploitation of supercomputing at national laboratories. However, the Panel finds that this position of leadership is seriously undermined by the lack of broad scale exploration, outside of a few national laboratories, of the scientific and engineering opportunities offered by supercomputing, and by a slowdown in the introduction of new generations of supercomputers. This threat becomes real in light of the major thrust in advanced supercomputer design that is being mounted by the Japanese Government and industry, and by vigorous governmental programs in the United Kingdom, West Germany, France, and Japan to make supercomputers available and easily accessible to their research and technological communities.

American preeminence in large scale computing has been a result of the confluence of three factors: the vitality of the U.S. computer industry, the far-sighted policies of the Federal government, and the leadership of scientists and engineers from universities and government laboratories. The Atomic Energy Commission, on the urging of John von Neumann, initiated the use of large scale computation in research and weapons design; NASA, prodded by Hans Mark, advanced the use of supercomputing in its scientific programs. American universities and government laboratories conducted the research that formed the basis for constructing and applying computers, trained the needed scientific and engineering personnel, and made computers and computing an essential tool in scientific and engineering research.

The Federal government vigorously implemented policies that supported these efforts, granted generous funds for computation, and, through its role as the major purchaser of scientific computers, provided the incentives and insured the market for these unique machines. Forward-looking corporations exploited the scientific and engineering opportunities, developed an advanced industrial technology, and created this most vital component of the American economy.

During the 1970's the Federal government retreated from its support of large scale computing in universities. The NSF program to provide and expand university computing facilities for scientific and engineering research was terminated in 1972; at about the same time IBM discontinued its generous discounts for the purchase of computing equipment by academic institutions as a result of pressures from the Justice Department and competitors. Since then large scale university computing facilities have withered while the action shifted to national laboratories and to industrial users. The most advanced scientific computer of the early seventies, the CDC 7600, was not installed on a single American campus, although it was available to researchers at several foreign universities and research institutes.

This continues today. With the exception of two universities and a few government laboratories, either dedicated to special tasks or specific disciplines, universities and government research installations lack supercomputers.[9]

Within the Government, fully integrated supercomputer facilities are found exclusively at dedicated national laboratories such as Los Alamos National Laboratory (LANL) and Lawrence Livermore National Laboratory (LLNL) in support of weapons and fusion programs; and NASA installations, the Geophysical Fluid Dynamics Laboratory, and the National Center for Atmospheric research (NCAR) in support of aerospace, oceanographic, and atmospheric research programs. The National Magnetic Fusion Energy Computer Center (NMFECC) of LLNL is accessible in interactive mode to researchers at remote locations by a high speed network.[10] On the other hand, the National Bureau of Standards and most DOD laboratories do not have supercomputers and far too few universities have the specialized computational equipment needed for scientific computing (e.g., array processors). As a result of limited access to supercomputing facilities by the broad research community, significant research opportunities[11] have been missed, and the younger generation of researchers is inadequately trained in large scale computing.

---

[9] See, in Appendix I, Partial Inventory and Announced Orders of Class VI Machines.

[10] See Appendix I for a description of some of these facilities.

[11] See, for example, the "Prospectus for Computational Physics," [2], "Future Trends in Condensed Matter Theory and the Role and Support of Computing" [4], "Report by the Subcommittee on Computational Capabilities for Nuclear Theory," [28], and "An Assessment of Computational Resources Required for Ocean Circulation Modeling," [35].

The need for access to large scale computational facilities has become so critical[12] that several universities, assuming significant financial risks, have felt it essential to acquire supercomputers. Several more are seriously considering doing so, and others are in the process of forming consortia for this purpose. Some of these endeavors have applied for Federal funding without which they may have financial difficulties. Other groups are pressing funding agencies to expand or replicate highly successful facilities, such as those at NCAR, NMFECC, and NASA-Ames, at universities or at national laboratories. Class VI scientific remote computing services are available from a few commercial service bureaus, but neither the academic nor the government research communities make extensive use of this resource. This seems due to a combination of lack of funds for computing services, the perceived high cost associated with these services, and a lack of sophisticated high-speed networking facilities. It is an indication of the absence of a national plan that a substantial number of leading scientists are clamoring for access to supercomputers at the same time that some supercomputing facilities are underutilized.[13]

A supercomputer is a general purpose scientific instrument serving a broad and diverse base of users. The decline of supercomputing at universities is analogous to the decline of instrumentation; neither large scale computing nor instrumentation can be sustained in a stable manner through funding of individual research grants, where their costs must compete with that of scientific personnel.[14]

The findings of the Panel regarding the development of supercomputers are as alarming as the findings on their access and availability. The U.S. supercomputing market is, at this time, dominated by Cray Research (CRAY-1) and Control Data Corporation (CYBER 205).[15] The Japanese vendors, Hitachi (S-210/20) and Fujitsu (VP-200), have announced the delivery of supercomputers in the near future, and these machines appear to be comparable to the available American systems. The Japanese are striving to become serious competitors of domestic manufacturers, and U.S. dominance of the supercomputer market may soon be a thing of the past. The Japanese Government-sponsored National Super Computer Project[16] is aimed at the development, by 1989, of a machine one thousand times faster than current machines. There is no comparable technical program in the U.S.. The Panel notes that in the case of the NASA Numerical Aerodynamic Simulator, a very high performance supercomputer, no acceptable proposals for its development were received. Neither of the two competing vendors could assure NASA that they would meet the performance requirements. Rather than developing new products, the vendors attempted to fit all NASA requirements to their existing product line.

---

[12] See, in Appendix III, The Supercomputer Famine in American Universities, by L. L. Smarr.

[13] Supercomputer cycles are available at the University of Minnesota CRAY-1 and at Colorado State University CYBER 205.

[14] See, in Appendix III, the paper by R. G. Gillespie. Most supercomputer facilities are funded directly by the Federal Government.

[15] See, in Appendix I, Partial Inventory and Announced Orders of Class VI Machines.

[16] See, in Appendix III, Japan's Initiatives in Large Scale Computing, by L. Lee.

Upon review of previous studies, the Panel also finds that the power of current and projected supercomputers is insufficient to meet existing needs in science, engineering, and technology, both military and civilian.[17] Research at universities and in the computer industry has indicated that future generations of very high performance computer systems may have parallel architectures radically different from the conceptually sequential architectures of today's supercomputers. There are many candidate architectures that must be evaluated before commercial feasibility can be established.[18] Simultaneously, the rapid and continuing advance of microelectronic technology makes it feasible to build such parallel machines. There is also a need for improvement of component performance.

The Panel believes that under current conditions there is little likelihood that the U.S. will lead in the development and application of this new generation of machines.[19] Factors inhibiting the necessary research and advanced development are the length and expense of the development cycle for a new computer architecture, and the uncertainty of the market place. Very high performance computing is a case where maximizing short-term return on capital does not reflect the national security or the long-term national economic interest. The Japanese thrust in this area, through its public funding, acknowledges this reality.

The Panel estimates[20] that the present annual investment in basic research on algorithms, software, and architecture is between 5 and 10 million dollars, while the annual advanced development expenditures for supercomputers (beyond Class 6 machines) are between 20 and 40 million dollars. This is contrasted with the development cost for a new high-speed conventional architecture system of approximately 150 million dollars, as well as the estimated 200 million dollars national superspeed computer project in Japan. The panel considers current levels of United States investments insufficient to maintain leadership in supercomputers.

The Panel believes that U.S. leadership in supercomputing is crucial for the advancement of science and technology, and therefore, for economic and national security.

---

[17] See, for some illustrative examples, Appendix II. Also notable is "Trends and Pacing Items in Computational Aerodynamics", by D.R. Chapman, [40].

[18] See, in Appendix III, a series of contributions to the Workshop. In particular, the papers by Dennis, Gajski, et al., Ris, and Fernbach; also the report of Group 3 of the Lax Panel in the Supplement.

[19] See, in Appendix III, Why the U.S. Government Should Support Research on and Development of Supercomputers, by B. Buzbee.

[20] Members of the Panel and of the Organizing Committee have conducted a survey to estimate the current total national investment in research and advanced development for supercomputer and supercomputing, both public and private. This survey included research and development costs but excluded funding for the acquisition, maintenance, and operation of supercomputer facilities.

Recommendations

The Panel recommends that the present needs and challenges to U. S. leadership in scientific computing and supercomputer technology be addressed with high priority. To this end, the Panel has set forth the background for planning and policy issues, outlined some options, and noted that current total funding in this area is insufficient to maintain the Nation's leadership in large scale computing. The Panel has avoided recommendations of a programmatic and organizational nature; these, and their implementation, are best left to the appropriate government agencies. These agencies must work together to respond to the issues raised and put together a detailed coherent program whose components are responsive to their individual missions. The program plan should contain a clear statement of goals, directions, and roles for the academic, industrial, and Federal government segments; responsibilities of the participating Federal agencies; and funding required.

The Panel recommends that a long-term National Program on Large Scale Computing should be initiated immediately, with the participation of the appropriate Federal agencies, the universities, and industry. The goals of this National Program should be:

1. Increased access for the scientific and engineering research community through high bandwidth networks to adequate and regularly updated supercomputing facilities and experimental computers;
2. Increased research in computational mathematics, software, and algorithms necessary to the effective and efficient use of supercomputer systems;
3. Training of personnel in scientific and engineering computing; and
4. Research and development basic to the design and implementation of new supercomputer systems of substantially increased capability and capacity beyond that likely to arise from commercial sources.

This Program should be coordinated by an interagency policy committee consisting of representatives of the appropriate Federal agencies, including DOC, DOD, DOE, NASA, and NSF. A Large Scale Computing Advisory Panel, with representatives from Government, the universities, and industry, should be established to assist in the planning, implementation, and operation of the Program.

The Panel finds two points that require emphasis:

> As the few successful facilities amply demonstrate, solution of the problem of access to supercomputing facilities, on a national basis, is possible.

Secondly, the domestic computer industry must allocate its scarce research and development funds to meet all the commercial opportunities and competitive challenges. Supercomputing enjoys a priority within the computer industry. But this priority, which reflects competitive commercial conditions, does not reflect the entire national scientific and security interest. It is not reasonable to rely solely on industry's own initiatives and resources in this area.

Possible Approaches for the National Program

The Panel has received many suggestions[21] for carrying out the thrusts of the proposed National Program. We outline here those considered most promising.

1. Access: There appear to be three approaches to provide reliable and efficient access to supercomputers to the research and development community. Common to all these is the development of a nation-wide interdisciplinary network[22] through which users will have access to facilities. This network should connnect all supercomputer facilities (except those dedicated to very special tasks), including commercial supercomputing centers and experimental machines.

   - o The most expedient and perhaps least expensive way to provide supercomputer access to the broad range of scientific and engineering researchers is to enhance supercomputer capacity and staff at existing centers which have demonstrated sophisticated capabilities for providing large scale computing.
   - o Provide supercomputers to selected government laboratories without such facilities and make them available to the broad research and development community through networking. In addition, there should be sharing and enhancement of current supercomputer facilities located at universities and government laboratories.
   - o Establish additional regional centers at selected universities, interconnected with existing facilities at other universities and government laboratories.

   The existence of a national network would permit combinations of these nonexclusive options, as well as the appropriate use of commercial services. The mechanisms for funding these facilities and allocating access should be carefully studied.[23]

---

[21] see position papers in Appendix III.

[22] The NMFECC network, described in Appendix I, is the example repeatedly mentioned for emulation because of its high bandwidth. The ARPANET network is often mentioned because of its interdisciplinary nature.

[23] See the position paper by R. G. Gillespie in Appendix III.

The above recommendations are related to the problems of access to general purpose supercomputer facilities. It should be noted, however, that there are scientific and engineering problems that can perhaps be better and more economically attacked by specialized supercomputing facilities and by sophisticated array processors.[24] The Panel recommends that, as part of the National Program, significant emphasis be placed on providing this specialized equipment to the research community. Finding the proper balance between investments on these two types of facilities requires a careful analysis, at the multidisciplinary and interagency level.

2. Research in Software and Algorithms

Today's supercomputers are a major departure from traditional sequential machines. Future significant improvements may have to come from architectures embodying parallel processing elements - perhaps several thousands of processors. In order to exploit today's vector processors and future parallel processors, entirely new algorithms must be conceived. Research in languages, algorithms, and numerical analysis will be crucial in learning to exploit these new architectures fully. The contributions of numerical analysis, computational mathematics, and algorithm design to the practice of large scale computing is as important as the development of a new generation machines.

3. Training

Another important component of this National Program is the development of an imaginative and skilled user community in supercomputing. There is a considerable shortage of appropriately trained personnel and of training opportunities in this area. Forms of institutional encouragement, such as NASA's special fellowships in the area of numerical fluid mechanics, special summer schools, and special allocation of access time to supercomputers for those projects that involve graduate students, should be considered. Some of the more mathematical aspects of these activities can be accomplished independently of the machines on which the actual calculations are done; however, the true integration of methods and their implementation cannot be done without access to supercomputers. The nature of the machine architecture has a very profound effect on the numerical methods, on the algorithms, and of course on the software. Thus, while being trained, students must have access to state-of-the-art computers. Today such training is virtually nonexistent; yet the skills gained from such training are essential to science and engineering.

---

[24] Some of these problems arise in a number of areas associated with experimental physics. See, in Appendix III, the letter from A. E. Brenner, Jr. See, also in the same Appendix, the position paper by K. Wilson on the use of array processors.

4. Research and Development for New Supercomputers

There are serious component and architectural problems that must be solved as part of the development of future generations of supercomputers. The unique strengths of industry, universities, and Government laboratories should be brought together for this purpose. A group of panelists from this workshop, with the aid of a number of experts from industry and universities, has produced a report[25] which describes one such program.

Since a great deal of careful analysis and detailed planning is required before the proposed National Program can be implemented, the Panel urges that its recommendations be acted upon as soon as possible.

---

[25] See, in the Supplement, A Program for Development of Very High Performing Computer Systems, by J. C. Browne and J. T. Schwartz.

A NATIONAL COMPUTING ENVIRONMENT
FOR
ACADEMIC RESEARCH

Prepared Under the Direction of
Marcel Bardon

by the NSF Working Group on Computers for Research
Kent K. Curtis, Chairman

National Science Foundation July, 1983

The views, budget projections and recommendations expressed herein are those of the participants in this study and do not necessarily imply endorsement of the National Science Foundation.

Members of the NSF Working Group on Computers for Research:

Darrel D. Aufenkamp, Mathematical and Physical Sciences
John W. Connolly, Materials Research
*Kent K. Curtis, Mathematical and Computer Sciences
Jerome S. Daen, Office of Planning and Resources Management
Thelma E. Estrin, Electrical, Computer, and Systems Engineering
Stephen Gould, Policy Research and Analysis
Richard A. Isaacson, Physics
Arthur Kowalsky, Physiology, Cellular and Molecular Biology
Lawrence A. Lee, Atmospheric Sciences
Edward A. McCullough, Jr., Chemistry
Michael J. McGill, Information Science and Technology
Jane T. Stutsman, Office of Planning and Resources Management
Alvin I. Thaler, Mathematical and Computer Sciences

*Chairman

Participants in the NSF Workshop on Computers for Research, May 23, 1983:

Dr. Alexandre Chorin, University of California, Berkeley
Dr. Arthur J. Freeman, Northwestern University
Dr. Freeman Gilbert, Scripps Institution of Oceanography
Dr. Walter Goad, Los Alamos National Laboratories
Dr. Richard J. Herrnstein, Harvard University
Dr. Franklin F. Kuo, SRI International
Dr. Cecil Leith, Senior Scientist, National Center for Atmospheric Research
Dr. Jacob Maizel, National Institute of Child Health and Human Development
Dr. Daniel P. Siewiorek, Carnegie Mellon University
Dr. Larry Smarr, University of Illinois
Dr. Larry Snyder, Purdue University
Dr. Frank H. Stillinger, Bell Laboratories
Mr. Joe B. Wyatt, Chancellor, Vanderbilt University

Participants in Special Meetings:

Dr. James F. Decker, Department of Energy
Dr. Robert E. Kahn, Defense Advanced Research Projects Agency
Dr. Kenneth G. Wilson, Cornell University

Executive Summary

Computing facilities have a decisive effect on the kind of research which is done by academic scientists and engineers. During the 1950's and 1960's, the Government encouraged the growth of computing in research, research methods were transformed in discipline after discipline, and the U.S. enjoyed a large, ever-widening lead in quantitative research and modelling complex phenomena. In the 1970s government support slackened and academic computing facilities no longer kept pace with advancing technology. During this time, advances in computer technology introduced minicomputers, microcomputers, automated measuring instruments and a variety of other computer dependent tools into academic laboratories; but academic researchers fell behind their industrial counterparts in obtaining and using the rapidly advancing computer technology. American universities installed no supercomputers and although two disciplines, atmospheric sciences and magnetic fusion energy sciences, established national computing facilities for their research, only a few scientists in other disciplines could gain access to supercomputers anywhere. The effect on academic science was not immediately apparent but scientists gradually realized that lack of opportunity to use sophisticated computing technology directs academic research away from those areas which demand it. This leaves vital fields unexplored and impoverishes American science. The lack of opportunity to properly train graduate students is of particular concern.

There are many organizations, both private and public, active in computer related research and education. For example, during the last few years NSF and other research supporting agencies have placed special emphasis upon improving academic research instrumentation. Some of these funds have been used for small computers. Also, computer manufacturers and other private sources are increasing their contributions to academic science and education, especially for computer based activities. Four universities have installed supercomputers with assistance from CDC, Cray Research and Denelcor but in general, support remains inadequate in all disciplines. In those disciplines that attempt to fund computing through individual research grants, access to large scale computing remains minimal.

In December, 1982, the National Science Foundation issued the report of the Panel on Large Scale Computing in Science and Engineering, Peter D. Lax, Chairman, [1] which urged the Federal government to take measures promptly to stem the erosion of American leadership. In its report, the Lax Panel stated:

> "Current and feasible supercomputers are extremely powerful scientific and engineering tools. They permit the solution of previously intractable problems, and motivate scientists and engineers to explore and formulate new areas of investigation. They will surely find significant applications not yet imagined. For these reasons, the Panel believes that it is in the national interest that access to constantly updated supercomputing facilities be provided to scientific and engineering researchers, and that a large and imaginative user community be trained in their uses and capabilities."

In April, 1983, NSF organized an internal working group on computers for research to help develop an appropriate course of action for the Foundation to meet the computing needs of academic science and engineering. The members of this group consulted with scientists in their respective disciplines and drew upon their own experience as program directors to estimate the needs in their areas of competence. On May 23, 1983, NSF held a workshop with thirteen scientists from diverse disciplines to define the needs better, especially for large-scale computing and networking. The working group also consulted many earlier studies, met with NSF advisory committees and individual scientists, and drew upon other advice from the scientific communities. This led to six principal conclusions:

- The attitude toward computers for research marks a generation gap in the community of scientists. Older scientists may use, or have their graduate students use, computers for a variety of measurement and analysis tasks but consider the role and importance of computers to be limited. Younger scientists, who grew up with computers as part of their culture, view them quite differently.

- A large gap exists between need and available support for minicomputers, attached processors, workstations, high precision graphics, local area networks and other local facilities required as part of the researchers' and graduate students' daily working environments.

- There is an immediate need to make supercomputers more available to academic scientists and engineers. The actual magnitude of the requirement is difficult to assess because lack of opportunity has artificially depressed the level of activity but an estimate of the near term need is for ten systems including an advanced vector computer at the National Center for Atmospheric Research.

- Computer networks are necessary to link scientists, wherever they are, to large-scale computing resources and to each other. Development and use of such networks should be encouraged.

- Science has passed a watershed in using computers for research. Computers are no longer just tools for measurement and analysis but large computers, in particular, have become the means for making new discoveries. Very important lines of inquiry exist which require computing capability 100 to 1000 times greater than any now in prospect. Academic research in computer architecture, computational mathematics, algorithms, and software for parallel computers should be encouraged to increase computing capability.

- There is competition among computer manufacturers in education and research. This is resulting in favorable terms for educational institutions to acquire computer equipment and offers NSF an opportunity to leverage its funds.

One other important observation emerged from the Working Group's studies, also. Academic institutions are on the threshhold of a personal computer revolution, led by a few institutions such as MIT, Brown and Carnegie-Mellon and assisted by some computer vendors, which is putting computing power comparable to large computers of the 1960's in the hands of every student. Although this has deep significance, especially for undergraduate and precollege education, the Working Group did not attempt an analysis of this phenomenon but focussed its attention on research and graduate education.

To meet the needs of U.S. academic research and graduate student training it is recommended that NSF should:

1. Develop and coordinate programs in this area with other federal agencies and with private performers and supporters of computer related research and education such as the computer manufacturers, the Semiconductor Research Consortium, Inc. (SRC), the Microelectronics and Computer Technology Corporation (MCC), and other private research performers.

2. Greatly increase its support for local scientific computing facilities, including:

   - workstations for individuals or small groups of scientists, e.g. microprocessors, graphics, data storage capability, terminals, and printers;

   - computer systems for groups of researchers and students;

   - specialized facilities such as array processors, symbolic processors, experimental computer systems, and scientific data acquisition, storage, and dissemination centers;

   - local area networks of computers and workstations;

3. Encourage proposals to provide supercomputer services for academic research and science education and be prepared to support 10 supercomputer systems in the next three years;

4. Support networks linking universities and laboratories with each other and with supercomputer centers to provide access to facilities, file transfer capability and scientific communication;

5. Establish an advisory committee to advise the Foundation on policies and mechanisms for supporting computer services and networks and to provide oversight;

6. Support a vigorous program of academic research and training in the areas of advanced computer systems design, including building and evaluating experimental computer systems, and in computational mathematics, software and algorithms necessary to the effective and efficient use of computer systems.

The cost of this program is estimated to be:

	FY 84	FY 85	FY 86
Local Facilities	$45M	$91M	$107M
Supercomputers	$15M	$70M	$110M
Networks	$2M	$7M	$11M
Advanced Computer Systems and Computational Mathematics	$8M	$20M	$33M
Total	$70M	$188M	$261M

These sums are realistic for a near term national effort in academic science and engineering but even these investments would not provide computing facilities for academic researchers comparable to those at a typical national or industrial research laboratory. Boldness of execution will be critical to success in reasserting American leadership.

I. THE ROLE OF MODERN COMPUTING IN SCIENTIFIC AND ENGINEERING RESEARCH WITH SPECIAL CONCERN FOR LARGE SCALE COMPUTATION

Background

W. O. Baker writing in 1970 on the topic of "Computers as Information Processing machines in Modern Science"(1) observed:

> "Science, the vast producer and consumer of knowledge, has until lately held mostly to handwork for helping headwork. Of course, there have been mechanical aids to printing and publishing the record of scientific advance and more recently a host of ways of sensing and recording the measurements of scientific experiments by instruments. But in between the original observation and its insertion into the great corpus of scientific record there has persisted a wide, in some ways a primitive, personalized domain. Of course, this gap in mechanization does happen to contain a most awesome function of mankind--the ability to think. Now, suddenly, digital machines have come into this realm of the scientist, supplementing mysterious talents to perceive and to recognize patterns, to organize data, and, above all, to envision in the mind's eye new concepts and understanding."

Baker continued:

> "Likewise, the applications of science for the benefit of man through industry and government increasingly depend on the conception and design of large interacting systems, in which many of the elegant uses of computers that we have sampled must be combined in the most ingenious ways. Far from seeing a plateau or decline in the rate of progress in a golden age of science and technology, I see boundless opportunities through the use of logic machine systems to resolve the puzzles of cosmology, of life, and of the society of man."

Well over a decade later, as Baker foresaw, computing permeates every field of scientific endeavor in increasingly profound ways. Kenneth Wilson writing in 1982 on the "Role and Needs for Large Scale Computing for Theoretical Physics"(2) stated:

---

(1) Baker, W.O. "Computers as Information-Processing Machines in Modern Science," Daedalus, Fall, 1970 pp. 1088-1120

(2) Report of the Panel on Large Scale Computing in Science and Engineering, Peter D. Lax, Chairman, December 26, 1982, Appendix II pp. 16-26

> "In conclusion, it is normal and reasonable that the computing demands for scientific simulation are prodigious, even by comparison with today's most powerful supercomputers. In fact, the whole area of scientific simulation in support of modern theoretical physics has barely been scratched; providing full access to theorists on the best of today's computers, can only be viewed as the beginning of many further stages of providing new levels of computing capability in support of ever more complex simulations. The complexity in simulation that can be handled numerically must keep pace with the growth in complexity that theorists are able to handle conceptually."

The use of modern computers in scientific and engineering research spans a period of only three decades, but it has been a dynamic period in which profound changes have taken place--and clearly the end is not in sight.

The period of the mid-50's through the 60's was one of significant change for research conducted in the academic environment; access to new tools--computers--was achieved. This was a period of the rise of campus computing centers in support of research and education. Institutional funds provided the bulk of the costs in achieving this objective but government support (and occasionally industrial support) was often a deciding factor in an institution's decision to establish or upgrade its computing facilities.

Although most institutions chose to develop computing facilities on an institutional basis, others were developing shared facilities--one notable example being the Triangle Universities Computer Center (TUCC) involving cooperation among the University of North Carolina, North Carolina State University, and Duke University. One unique facility in the 60's, which provided services to researchers at many institutions, was the Western Data Processing Center at UCLA, a "large scale" center in its time established with support from IBM. IBM also provided support to UCLA in the late 60's together with support from NSF for a major campus computer facility featuring the IBM 360/91 computer, a "supercomputer" of this period. Similar large scale computing facilities featuring Control Data Corporation, IBM, or UNIVAC equipment and support were established at several other universities, also. A characteristic of these facilities was that they were developed for multidisciplinary use--providing computing services to researchers without regard to discipline as well as for educational use and often for administrative applications.

A different approach which provided extensive computing facilities for a single discipline or a group of closely related research activities was also started during the 1950s and 1960s. The U.S. Atomic Energy Commission encouraged extensive use of computers for management of large scale experimental facilities and data reduction in high energy physics. This was seminal work which greatly developed the application of computer technology to research. NSF organized the National Center for Atmospheric Research (NCAR) to be an intellectual center for major research efforts and a focal point for developing and housing the expensive and sophisticated research facilities required for effective progress in atmospheric research. The scientific computing facility was created within NCAR shortly thereafter. Computing resources were, and still are, allocated free of charge to NSF-supported atmospheric and oceanographic scientists in the universities and at NCAR.

Today, the computational power provided by NCAR is needed by atmospheric scientists and physical oceanographers to pursue the studies required for the advancement of the sciences. NCAR provides computing capability of a magnitude sufficient to support computations for extensive modeling and data analysis as well as storage and archive capacity for major data collections. The emphasis is on the best available, high-performance computers and support services tuned to this research.

The Scientific Computing Division at NCAR serves approximately 850 scientists through two CRAY-1A computers with IBM 4341 front-end systems, and a full set of graphics capabilities. Services provided to users from the staff of 100 individuals include: 24-hour per day operation throughout the year, consultation on systems, applications and numerical analysis, and large software libraries. Meteorological data bases are also archived and distributed.

The large computer facilities at the Department of Energy and NASA national laboratories also represent disciplinary-oriented facilities within the mission of their respective sponsoring agencies.

The Health Sciences Computing Facility at UCLA, supported by the National Institutes of Health, was a major computational research and services center in the 60's and early 70's. The Biomedical Statistical (BMD) Software packages, used extensively by researchers in many disciplines world-wide, originated at this center.

The horizons for the role of computing in scientific and engineering research broadened extensively with developments which began in the 60's to improve the access to computers. Remote access coupled with time sharing and interactive systems permitted users physically remote from the computer to make concurrent efficient use of a computer facility not only for creating and debugging programs on-line, but for production runs and subsequent analyses of data, often in an interactive mode. The sharing of computers does not in itself require remote access service; individuals can and, in many cases, do go physically to the site of the computer in order to carry out their research computing. Remote access, however, through specially designed computer-communication facilities has led in a natural way to "networks" to improve computer access and user services for resource sharing. For example, remote users of the NCAR computing system access that system through the national communications network, UNINET. This is a low cost, low error rate, packet switching service with 177 nodes scattered throughout the Nation. In order for NCAR to provide the same level of computing capability as is provided to the 300 or so users at the Center, a plan of transmitting very large blocks of output data across the country at high speeds is being developed.

The period of the early 70's was a fertile one for coupling computers and research. Mini-computers became available and it was during this period that imaginative scientists in many disciplines began collaborating with computer scientists to make advances in research methodology. Achievements were many: Laboratory instrumentation became much more sophisiticated as computers became an integral part of the system permitting monitoring and feedback on a real-time basis and coupling the researcher much more directly with the experiment. There was much exploration of the use of computer graphics. These sophisticated laboratory instruments were often in turn part of a local hierarchy, a network of computer facilities for more extensive data storage and analysis. A major computer research center for economics and management science was established at the National Bureau of Economic Research for advancing research methodology in those disciplines. A special project centered at Argonne National Laboratory (also supported by NSF) involved the collaboration of researchers on an international basis to create "certified" high quality mathematical software for a wide range of computers. This software and its successors are now widely used and have been incorporated as well in commercial software products.

In the early 70's some universities had the "supercomputers" of the period; however, by the middle 70's there were no supercomputers in universities. NSF support for campus computing centers was discontinued in FY 1972 and the research community turned to work which could be done with facilities available; those few academics who persisted in large scale computing invested large amounts of time and energy seeking time on facilities at Los Alamos National Laboratory (LANL), Lawrence Livermore National Laboratory (LLNL) or in Europe.

Other changes were taking place as well: computing became an essential component of research in more and more disciplines and, technological developments which had led to the appearance of minicomputers later brought microcomputers and then supermicro's, VAX's and array processors. Consequently, academic departments or even individual research groups could, in effect, acquire their own computer centers. These developments coupled with a common practice of Federal granting agencies not to support the full cost of computer time led to a decentralization of campus computing. In many places, campus centers fell into disrepute as providers of services for research because of operating policies favoring small jobs, educational or administrative uses. During that same period, lacking federal support for research instrumentation, academic facilities fell behind industrial laboratories in computing resources for research including minicomputers, work stations, graphics, etc. as well as supercomputers.

This image of undercapitalization of academic computing in the 1970s is confirmed by the data gathered in academic year 1976-1977 by John Hamblen and Carolyn Landis for EDUCOM [36]. They discovered that the total expenditures for computing in higher education was $991M in 1976-1977. Of that total only 7 percent came from federal funds. 82 percent came from institutional funds and 11 percent from other sources. At major universities the proportion of federal funds was slightly larger at 10 percent. By contrast, in 1966-1967 the proportion of computing expenditures at major universities provided by federal funds was 34 percent. For higher education as as whole, it was 28 percent. Clearly, federal support did not keep pace with the expansion of need for academic computing.

It was in this milieu that the National Resource for Computational Chemistry (NRCC) was established. In their analysis of the Phase I operation of NRCC [35], the Ad Hoc Review Committee pointed out that the computational resources available through NRCC had been far too limited to have major impact. The capability which was provided (partial use of a CDC 7600) was not state-of-the-art; access to NRCC was never trouble free, and the available computing time was much too small. Furthermore, powerful minicomputers had recently begun to appear on the scene, and little of the research conducted at the very limited facilities of NRCC could not have been accomplished at local facilities, probably at lower cost and higher human efficiency. Since NRCC never did offer an advanced capability on the level of, say, NCAR, the failure of NRCC to produce any frontier research requiring such capability was a foregone conclusion. The chemistry community did not unite behind the vision of NRCC and its potential was never realized.

On the other hand, the DOE Magnetic Fusion Energy Computer Center (MFECC), dedicated to fusion physics calculations, was also established in the mid-1970's following recommendations of an Ad Hoc Panel on the Application of Computers to Controlled Thermonuclear Research. This center started operations in late 1975 with a CDC 6600 followed shortly by a CDC 7600. A review of requirements in early 1976 resulted in the procurement of a CRAY I within two years and a second CRAY I in 1981. The long term commitment to supercomputers in this area coupled with the development of a communication system for remote users and a strong user community has led to a very successful utilization of this center.

The failure of NRCC and the success of MFECC depended critically on the non-availability or availability of supercomputers and the timidity or boldness of planning for each; the essential role of computing, and that of large scale computing in particular, for the disciplines concerned was not at issue.

In recent years there have been many studies, culminating in the Lax Report [1], which call attention to a need for change in the policies of the government vis-a-vis computing for academic research. In response to this concern, the National Science Foundation established an internal working group on computers for research in April, 1983, to help develop a plan of action for NSF. The members of this group consulted with scientists in their respective disciplines and drew upon their own experience as program directors to estimate the needs in their areas of research. On May 23, 1983, NSF held a workshop with thirteen scientists from diverse disciplines to define the needs better, especially for large-scale computing and networking. The working group also consulted with NSF advisory committees and individual scientists and drew upon other advice and publications from the scientific communities. This report presents a summary of the working group's conclusions and recommendations. A brief overview of large scale computing facilities available for scientific and engineering research is provided in Part I.A. This overview includes national laboratories, universities and selected commercial service centers. More detailed information about the individual facilities appears in the Appendix. Part I.B. contains a critique of current facilities and support programs. In Part I.C., the type and amount of additional computing services needed is given. Section II. of the report analyzes the responsibilities of NSF for academic computing which result from the Foundation's charter and its relationship to other federal agencies. Section III. contains recommendations for action by NSF and Section IV. provides budget estimates for carrying out those recommendations. A partial bibliography of relevant reports and other materials is attached, beginning on page 45.

1. A. Summary of Current Uses and Support of Large Scale Computing for Scientific and Engineering Research

In this part, the supercomputer facilities of National Laboratories which offer access, in part, to academic researchers (within the constraints of the missions for which the facilities were established) are indicated. Brief descriptions of their major computer systems are given, as well as the research community served, communication access for remote users, charges, and plans for future expansion. Information is provided as well on three existing university supercomputer systems and three commercial service centers. A description is also provided of an ambitious plan to increase local computing capability at one major university, Cornell.

FACILITY	MAJOR SYSTEM	RESEARCH USERS	CHARGES
NCAR	2 CRAY 1-A's	Atmospheric and Ocean Sciences	No charge to NSF users; $2,200 per prime CPU hour for others
MFECC	2 CRAY 1 1 CDC 7600	Magnetic Fusion Energy Community	No charge
LANL	(Open): 1 CRAY, 3 CDC 7600's, 1 CYBER 825	Government Agencies, Labs and Non-profit Institutions	$636 per prime CPU hour
NASA-Ames	CRAY 1-S CDC 7600	Computational Fluid and Aerodynamics	No charges to NASA grantees; $2,000 per CPU hour for CRAY
NASA-Goddard	CYBER 205	NASA-funded and NASA-project related	No charge to NASA grantees; $1,000 per CPU hour for others
NASA-Langley	CYBER 203	NASA and NASA-funded scientists	$1,300 per CPU hour
NASA-Lewis	CRAY 1-S	Principally Aerodynamics-related	No charge for NASA-supported
Colorado State University	CYBER 205	University users and others	$2,000 per hour for university users

FACILITY	MAJOR SYSTEM	RESEARCH USERS	CHARGES
Minnesota	CRAY 1-A	University users and others	$2,000 per hour for University and non-profit; $3,000 per hour for others. Volume discount available.
Purdue	CYBER 205	University users	Not yet determined
Boeing	CRAY 1-S		
CYBERNET	CYBER 205	All	Variable
U. I. S.	CRAY 1-S	High-tech., insurance companies, etc.	Variable
Cornell	IBM 370-168 Several 4341's FPS AP-190L FPS 164	50-100 Cornell scientific users	Typical user rates: $20-40 per CPU hour (non-profit) for array processors; $500-$1,500 per hour on IBM hosts

A more detailed description of each of these facilities is given in the Appendix. In particular, the description of the facilities at Cornell gives a picture of an ambitious program attempting to provide local computing facilties to research groups comparable with the capabilities supplied at major national facilities.

From this list, it is clear that except for academic researchers in certain restricted groups, e.g., atmospheric sciences or NASA supported, or a few who can form collaborations with researchers at the Los Alamos or Livermore National Laboratories, there is no access to supercomputer facilities without substantial funds. The unit cost of acquiring and operating supercomputer facilities is very high, on the order of $10 million. It has not proven possible to accumulate these costs through recharges to individual research grants or contracts because too many competing interests interfere with the concentration of funds for computing. At the level of the research supporting agencies, confronted by inadequate research budgets, support of graduate students, research scientists and small laboratory instrumentation exhausts the available funds. At the level of the individual recipient of research support, funds granted for computing are recognized as inadequate and often diverted to other purposes. This has discouraged the development of computer-based research. Apparently, extraordinary measures are required to assure access to supercomputers. Dr. Trivelpiece, DOE, noted in his testimony to the joint hearings of the Subcommittee on Energy Development and Applications and the Subcommittee on Energy Research and Production of the House of Representatives Committee on Science and Technology on Tuesday, June 7, 1983, that in the Department of Energy's experience one must allocate computer time, not money. Otherwise, sufficient access to supercomputer facilities for research cannot be assured. NSF's experience is consistent with DOE's. NCAR, which allocates free time on its supercomputer facilities to meritorious atmospheric sciences research has been very successful in serving the academic atmospheric science community. Few other NSF supported scientists have been able to gain access and then only with great expenditure of effort.

In closing this section, reference should be made to the growing use of supercomputers by industry. Oil companies have turned to supercomputers for assistance with exploration for new deposits of oil and gas reserves and for tertiary recovery of oil deposits. Manufacturers of advanced microelectronic circuits, involving hundreds of thousands of logical elements on a chip are using supercomputers to check circuit designs for completeness and consistency. Airplane manufacturers have used supercomputers for many years to assist with the analysis of aerodynamic flow around proposed designs and have found that computer analysis can greatly reduce the very expensive tunnel testing time required to perfect new designs. Even movie makers and television have turned to supercomputers to create the elaborate visual effects that have appeared on recent films. The private sector has offered a new and substantial market for supercomputers in the past few years and it is expected that these applications will grow in number as the computing power of supercomputers increases to permit more realistic simulation of complex, three dimensional, time dependent phenomena.

I. B. Critique of Current Facilities and Support Programs

There is much positive to be said about the computer facilities and support programs currently available--particularly the efforts of communities of users, whether by discipline or by organizational entity, to develop computer facilities, networks, and services for resource sharing, and the efforts of the commercial sector as well to provide high quality computer-communication links. Nevertheless, important scientific and engineering research problems are not being attacked at present for lack of computer facilities with sufficient capacity and capability and appropriate access (network or otherwise).

To turn specifically to the concern for large scale computation in scientific and engineering research, a closer examination of the present situation in regard to capacity, capability and access reveals many deficiencies. The DOD, DOE and NASA facilities are reserved for their respective contractors or at least for problems within the missions of these agencies. The NCAR computer facility is fully utilized by NCAR and the academic community in the atmospheric and related areas. University supercomputer facilities have been established only recently. User communities and the underlying financial bases are being developed. The latter is uncertain of accomplishment through individual research project grants. The machines are considered to be under-configured currently for many problems. This may also be a factor, in part, for their present underutilization. Furthermore, none of the existing university facilities is presently equipped to provide the services necessary to serve a large community of off-campus users. Aside from the hardware and software, this includes services such as remote consultation, training, etc.

The situation with respect to U. S. facilities for academic research is in marked contrast to that in some foreign countries, W. Germany, for example, where adequately configured machines are readily available to scientific researchers. Indeed, some U. S. researchers have actually traveled to foreign countries to perform large scale computations that they were unable to perform here because of lack of access to national laboratory facilities or lack of funds to purchase supercomputer services. The fact that this inefficient and costly method of performing research was consciously chosen indicates the inadequancy of the situation with regard to access to large scale computing facilities in this country. It also illustrates that the existing NSF mechanisms for support of computational science do not appear adequate to remedy this situation.

Computer networks are an important vehicle through which to improve the development and utilization of information and distributed computing resources. There are many existing networks, some of which (e.g., MFE) include nodes providing large scale computing capacity. However, like most current supercomputer facilities, access to these networks is restricted to special communities. These communities do form important subsets of the total scientific community, so the existing networks do provide a substantial base on which to build. Unfortunately, sufficient motivation for extensions of the necessary scope and flexibility is unlikely to be generated from within existing networks.

Although the barriers to such extensions are mostly political, organizational and economic rather than technological, they are still sufficiently large to inhibit unilateral action.

An assessment of large scale computing for scientific and engineering research is provided by the many studies and plans in the reports listed in the bibliography.

Thus, the vital role of computing in research, and especially the role of large-scale computation, is well established. Problems are now being addressed, as Professor Kenneth G. Wilson indicated, which require tens to hundreds of hours of computer time of computers operating at hundreds of millions of operations per second. In some problems, data in quantities measured in trillions of bits are analyzed using sophisticated computer graphics systems to augment human powers of analysis. Without computers these computations and analyses would be completely impossible to carry out.

While capacity(3) is one concern, a greater concern for the future is capability(4). Some persons describe these as brute force and turn around time respectively. At both the very low end (personal computers) and the very high end (supercomputers) capability is the limiting factor on the problems which can be solved. The elapsed time or scientists waiting time must be realistic to make a line of research feasible. The participants in the Foundation-sponsored workshop on Computers for Research on May 23, 1983, indicated several problems that they would like to solve but can not because of the capability of current computers. Among these were:

- (BIOLOGY) The simulation of a bi-layer membrane in contact with ions in water would permit tracing the paths determined by polypeptides for ion transport through the membrane. It is estimated that computational power of 1000 times a CRAY is needed for this fundamental problem of biological significance.
- (CHEMISTRY) Chemical reactions can be modeled using the basic equations of quantum mechanics. To model reactions involving larger, more interesting, molecules substantially more computational power is needed.

---

(3) Capacity is defined as a measure of the volume of work a specified configuration can accomplish in a unit of time.

(4) Capability is defined as a measure of the highest degree of complexity or largest size of one job that a computer or network can process in some specified time frame.

- (MATERIALS RESEARCH AND ELECTRICAL ENGINEERING) As dimensions get small in microelectronics, classical physics fails. Quantum mechanics is needed to model small dimension physics for semiconductor devices.

- (CHEMISTRY AND MATERIALS RESEARCH) To understand catalysis on a surface, extensive calculations beyond present capabilities are required.

- (ATMOSPHERIC AND OCEAN SCIENCES) Turbulence calculations require three dimensional simulations which can not be done with current computational capabilities.

- (ATMOSPHERIC AND OCEAN SCIENCES) Calculations for atmospheric modeling require finer resolution of grid points and longer time scales; modeling ocean basins would require about 100 times the current capability.

- (EARTH SCIENCES) Modeling dynamic magnetic systems,--for example, to determine why the Earth's magnetic field reverses--requires capability beyond that currently available.

- (ENGINEERING) The computer capability does not yet exist for effective non-numeric processing required for robotics and for knowledge-based systems.

- (ASTRONOMY AND PHYSICS) Astrophysical problems as modelling the processes by which stars are born require much greater computer capability than current supercomputers offer.

- (PSYCHOLOGY AND INFORMATION SCIENCES) Current computers are not sufficiently powerful to model a human's complete knowledge, even of one limited subject area.

- (ECONOMICS) Computer technology now limits the rigor and size of dynamic nonlinear economic modeling.

Strong research efforts in computer architecture, computational mathematics, algorithms and other aspects of computer research will be required to achieve the significantly increased computer capabilities required to address these future generations of research problems. A concerted effort is therefore needed to strengthen substantially academic research programs in computer systems--from which most significant advances have come--in this search for new computer architectures applicable to the scientific and engineering research problems of the future.

Supercomputers and large scale networks represent only one end of the spectrum of capability needed by computational researchers. Many, if not most, scientific computations can be performed with tools of lower capability. These are tools which should be locally accessible to a researcher: campus mainframes, departmental minicomputers, personal workstations, etc. Indeed, with the advent of special purpose processors, local facilities offer the most cost effective and humanly efficient means of increasing both access and capacity, within the limits of their capabilities . Furthermore, adequate local facilities are absolutely essential for the success of any distributed large scale computing system. Any examination of the overall needs of computational science must consider local facilities as an integral component of the total need.

Disciplines vary widely in the degree to which computer use has developed within the community. It can be seen that disciplines also vary widely in the mix of need for capacity across the spectrum of capability. Nevertheless, whatever the mix, the unmet need is substantial.

I.C. Unfilled Needs for Computer Support of Research

Many studies have documented the unmet needs for computer support of research in a qualitative fashion. Quantifying these needs is more difficult, especially with respect to large scale computation. To assist in this effort, various information sources were utilized. One of these was a workshop held on May 23, 1983. The participants, who are listed on page 2, were computationally oriented scientists drawn from the whole spectrum of disciplines, representing not only academics, but also industry and the national labs. They were asked to consider the unmet needs in four areas:

- large scale computing capacity
- local facilities
- networks
- new capability

Regarding the first, there is unanimous agreement that important science is not being done in this country due to lack of access to adequate supercomputer capacity. If an academic scientist does not belong to one of the specialized research communities with guaranteed access, it often requires heroic efforts to scrape together adequate access. For example, important astrophysical calculations essentially could not be performed in the U. S. because they did not fall into one of the areas (atmospheric sciences, magnetic fusion, etc.) covered by the specialized facilities, and the few university supercomputers were inadequately configured for the task. In this particular case, the researcher resorted to travelling to West Germany, where a collaboration with scientists having adequate supercomputer access allowed the calculations to be performed.

Needs for large scale computing capacity will grow rapidly as more disciplines acquire computer experience. These needs may arise in unexpected ways. For example, the biology community has established a nucleic acid sequence data bank at Los Alamos to make this vital information generally available to researchers. Unfortunately, as the data bank grows large even the most elementary manipulations of it, such as finding all entries matching a given sequence, consume several minutes on current supercomputers. When multiplied by the many demands, the needed supercomputer capacity greatly exceeds that available to the data bank.

The available evidence indicates an urgent need to make a capacity equal to at least two adequately configured supercomputers available to the general scientific community immediately. The best projection is that as many as ten will be required to have any major impact on the total unmet need across all areas of science.

To keep the number ten in perspective, it is worth noting that the magnetic fusion energy community alone currently has access to two supercomputers, and the Los Alamos National Laboratory has five.

Other types of large scale computing facilities are also important for various disciplines. The nucleic acid sequence data bank, cited above, is one example which is of interest to biology and chemistry. For computer research, and increasingly for other disciplines as they also engage in the design of special purpose computers, a service to fabricate custom-designed Very Large Scale Integrated (VLSI) circuits will be needed. The MOSIS facility, sponsored jointly by NSF and DARPA, is a prototype of such a service. Other specialized large scale facilities will become important as computers are used more in research.

A network appears to be the only logical vehicle for making this capacity accessible to the general scientific community. Experience with existing networks (e.g., MFE net) has been very positive. Workshop participants were in general agreement that a network is an essential component of any response to unmet needs in large scale computing. A network can also facilitate the interaction between scientists and promote the exchange of information, in addition to serving as a medium of communication between scientist and computer.

Regarding local facilities, there is agreement that they must be significantly strengthened, and that this must be undertaken for a balanced response to computational needs. Supercomputer access should be managed as a scarce resource. Indeed, any calculation which does not require supercomputer capability should be excluded from access to supercomputers and performed on other facilities.

Such calculations are most efficiently performed locally, and since such calculations will comprise the bulk of important scientific calculations, substantial local facility capacity is needed. Although funding patterns have changed radically in response to the shift toward decentralized local computing spawned by technological advances, funding levels have not kept pace with needs. All disciplines now experience unmet needs for local facilities.

The preceding discussion has dealt entirely with access to capacity. Scientific researchers also need access to machines of vastly higher capability. Factors of 1000 tend to be mentioned routinely. While factors of 10 are foreseeable with technological improvements, no one knows how to build a machine 1000 times faster than those currently available nor how to develop problem strategies or algorithms for them.

Thus, in contrast to the other unmet needs, the need for increased capability is the ultimate stumbling block. In all disciplines, progress in computational science will simply grind to a halt unless this impasse is resolved. Obviously, from a long term perspective this represents the most urgent need of all. Since it is also fundamentally a research problem, its resolution may be slow and painful; consequently it needs immediate attention.

II. The Role and Responsibilities of the National Science Foundation with Respect to Modern Scientific Computing

A. Introduction

The role and responsibilities of the National Science Foundation with respect to modern scientific computing must be defined within a larger context: the role and responsibilities of the Federal Government in research and development.

Substantial investments in basic research in all areas of science and engineering are widely recognized as essential for the long term health, prosperity and security of the Nation. It is also widely recognized that it is the responsibility of the Federal Government to provide the necessary investments in cases where there is insufficient incentive for the private sector to do so. This recognition is reflected by the increased attention now being focused on Government-wide support of basic research in general.

Within the overall Federal role, the National Science Foundation has the lead responsibility for basic research and science education conducted in academic institutions. By the National Science Foundation Act of 1950, as amended, the Foundation is authorized and directed to

> "...initiate and support basic scientific research and programs to strengthen scientific research potential and scientific education programs...in the mathematical, physical, medical, biological, engineering, social, and other sciences...
>
> and
>
> "...foster and support the development and use of computer and other scientific methods and technologies, primarily for research and education in the sciences..."

Thus, from both a general and a specific perspective, the Foundation has a mandate to assume a major responsibility with respect to computing for research and education in the sciences and engineering.

B. NSF as a Supporter of Access to Computer Capacity

Historically, the Foundation has interpreted this responsibility to include programs designed to promote access to and improve the capacity and capability of whatever computing facilities may be necessary for the solution of fundamental problems in basic research as well as education in the use of such facilities.

For example, from 1957 to 1972, the Foundation awarded grants to help initiate and expand computing facilities for education and research at colleges and universities. In that era, large mainframe computers in campus computing centers were the research instruments for the comparatively few computationally oriented scientists, including the practitioners of the then novel discipline of computer science. This program was instrumental in insuring that adequate computational facilities were available to these scientists. The program had a significant educational component as well. It played a major role in the development of a body of computer literate individuals in the U.S. more than a generation before any other country. The focus of the program was always stimulation rather than sustenance, and few institutions received more than one award. The program was finally discontinued when it was perceived that the providing of computer services had become sufficiently incorporated into the normal operation of major research universitities that special Foundation assistance was no longer required. The Foundation then reverted to supporting scientific computation by the distribution of funds for the purchase of computing services through individual research grants.

Another illustration is provided by the special computing facilities that have been created on occasion by the Foundation in response to specific perceived needs. The two notable examples are NCAR and NRCC. The former has, of course, been very successful in its mission, while the latter was not successful, for reasons mentioned previously. These two contrasting examples point out the importance of careful planning and highlight the danger of proceeding with inadequately configured facilities. A submarginal response to a specific need may be worse than no response at all.

The trend toward a more decentralized environment for scientific computing which was sparked by the minicomputer and accelerated by the microcomputer is generally a very positive direction. These devices offer a quantitative leap in both the access and capacity that can be made available to the scientific community for a given investment. The supermicro of today has the capability of the most powerful mini of a few years ago, while a modern mini with attached processor can put the equivalent of a CDC 7600 into the hands of an individual department or research group. Recognizing the cost-effectiveness of this trend, many areas of the Foundation have already been quite receptive to proposals requesting assistance in the purchase of minicomputers or microcomputers. This receptive posture should be extended to the Foundation as a whole. The unmet need is very great, as indicated in section I. Furthermore, minicomputers and microcomputers for scientific research have no fundamental distinction from other equivalent cost instruments for scientifc research routinely requested from the Foundation. The only difficulty with adopting such a posture is that it can only increase the competition for already scarce instrumentation resources.

C. NSF Policy vis-a-vis Large-Scale Computing

While increased support for the acquisition of minicomputers and microcomputers would go a long way toward alleviating many needs for increased access and capacity, such a policy can have no affect on unmet needs for access to computing systems of the highest capability. These systems are known generically as supercomputers. At any given time, only a small fraction of basic research problems might truly need supercomputer capabilities. However, in many disciplines, it is precisely these problems which are encountered at the frontiers of progress.

The current situation in the U.S. with regard to supercomputers is far from optimum on almost every count. Supercomputers are manufactured by only a few vendors. The total number of machines is quite small. None of them is easily accessible by a member of the general scientific research community, and most of them are not accessible by that community at all. In particular, only four supercomputers presently may be found at universities.

It has also become evident that although the U.S. has always been the leader in the development and use of supercomputers, that leadership is being challenged on a number of fronts.

The role and responsibilities of the National Science Foundation with respect to the special problem posed by large scale computing and supercomputers must also be defined within the overall Federal role. Several Federal agencies have already proposed attacks on various facets of this problem.

The Department of Defense, through DARPA, has established a program aimed at the development of strategic computers to meet specific DOD objectives. While the goal will not be machines for general purpose scientific computing, the program will undoubtedly provide a great stimulus for certain aspects of computer science and engineering.

NASA is planning to develop a new Numerical Aerodynamic Simulator (NAS). NAS will include a specialized large-scale computing system that will perform calculations required to design aircraft and significantly reduce wind tunnel and flight testing. In space research, NASA expects its computational workload to increase from one to ten billion operations per second. To meet this need, the Massively Parallel Processor has been developed and processing research will begin this summer. NASA has large-scale computational capacity at the NASA centers. These facilities are dedicated for use by NASA scientists and engineers, NASA contractors, and NASA-funded researchers.

The Department of Energy is investigating ways to make its own large scale computing facilities more accessible to qualified researchers who need such capabilities. They also are examining how the Federal Government as a whole can stimulate the commercial development of supercomputers. Vendors generally agree that the most important first steps are for the Government to maintain a "friendly customer" attitude and to provide industry with realistic projections of minimum government needs in the area. Since future needs in capability probably will require some radical departures from current designs, the "friendly customer" attitude might extend to experimental prototypes which show reasonable promise of achieving breakthroughs in capability.

In view of the Foundation's responsibility for basic research and education in science and engineering, its role with respect to large-scale computing and supercomputers should be to maintain strong programs designed to provide:

- support of basic research and science education in universities in computational mathematics and advanced computer concepts including software, algorithms, architecture, subsystems, and components;
- support of basic research in problems that require large scale computing across all areas of science and engineering;
- training of scientists and engineers in the concepts, design, and use of large scale computing systems;
- assured access for researchers and graduate students to large scale computing facilities.

If the words "large scale" are deleted, the preceding four statements provide an effective summary of the role of the National Science Foundation with respect to modern scientific computing in general. The Foundation recognizes that item four is crucial for the success of items two and three, and has assigned a high priority to resolving problems of access at all levels of capability.

These responsibilities of the Foundation with respect to scientific computing are not shared with any other agency. In particular, the problem of assured access to large scale computing for academic scientific research is not addressed at all by the DARPA strategic computer program, and only slightly by the NASA and DOE programs. Thus, the Foundation should assume the lead role in the resolution of this problem in accordance with its responsibilities.

Supercomputer systems probably will never again be commonplace on university campuses. The image of a supercomputer in every academic computing center appears quite unrealistic. The capital and operating costs of such a facility are simply too great relative to the average need that will be generated from within a single university.

Given that this situation is likely to persist into the foreseeable future, attempts to support large scale scientific computing solely by the distribution of funds through research grants probably are doomed to founder on the problem of access. The Foundation's experience has shown that academic institutions are often reluctant to spend grant funds for off-campus services, especially in computing. In addition, both NSF staff and researchers, when faced with budget restrictions, often sacrifice computer time budgets first as being most easily postponed. In brief, willingness on the part of the Foundation to provide funds for large scale computation in research grants is, in itself, insufficient to insure that scientifically pressing large scale computations can, in fact, be done. Insuring access to large scale computing facilities requires more drastic action ranging from contracting for large blocks of time with existing suppliers to creation of dedicated facilities. The Foundation has the responsibility to consider special solutions of this problem, including the possible role of large scale networking.

The fact that special solutions are required indicates that access to large scale computing for scientific research is a special resource. Inevitably, this means that it will be a scarce resource, and it should be treated as such. The goal of the Foundation should be to insure that a scientific researcher has access to precisely the capacity and capability needed for the particular problem at hand. Since only a small fraction of the time will this imply supercomputer capability, any policy which attacks only this aspect of the problem makes little sense, either scientifically or economically. There are substantial needs at all levels of accessibility, capacity, and capability, across all disciplines. All these needs must be addressed for maximum results to be obtained.

III. A Plan of Action for the NSF: Recommendations

A. Introduction

Section I. of this report has shown that the presence or absence of computing facilities has a decisive effect on the kind of research which is done. The failure of the U.S. to provide supercomputer facilities for academic research throughout the 1970's and early 1980's discouraged the development of U.S. science in the direction of large scale, complex modelling except in those disciplines such as atmospheric sciences, fusion research or aerodynamics which specifically provided supercomputers for their researchers. However, there are strong indications from the scientific community that urgent action is needed. For example, the initiative of Nobel prize-winner Professor Kenneth G. Wilson, Cornell University, whose current work has its principal focus on numerical modelling of physical phenomena and who is publicly urging greatly increased support for computer-based research [34], is a signal that we can no longer procrastinate. U.S. scientific and engineering research is being blunted by the lack of appropriate computing power.

During this same period, the U.S. has not provided the wide variety of computing services its academic scientists need to keep them in step with modern technology. Supercomputers are only one component of the collage; and minicomputers, attached processors, personal workstations, high resolution graphics, symbolic processors, local area networks, etc. are indispensable for modern research productivity. Industry, which recognized this for its own R&D, makes large investments each year for its scientists and engineers to supply computer equipment, software and support personnel. The federal government needs to do the same for academic research.

Section III. of this report presents recommendations for a plan of action for NSF. The Foundation should act on four levels:

a. increase support for local computing facilities;
b. support supercomputer services and access thereto for academic scientific and engineering research and education;
c. assist with the formation and use of appropriate computer communications networks;
d. support academic research in advanced computer systems design and computational mathematics to improve our computational capability for solving problems beyond the reach of todays supercomputers.

Each of these is treated in turn.

B. Coordination

As has been noted several times, there are many active participants in computer-related research at all levels. Several federal agencies are involved in or support computer-related research relevant to their missions. The computer manufacturers and some other companies such as Bell Laboratories, General Electric, Exxon, Xerox and General Motors have substantial in-house research programs and help support academic research and education. The SRC and MCC are supporting and conducting research, respectively, in relevant areas and several non-profit corporations such as RAND, Draper Laboratories, SRI International, and MITRE are research performers. Also, there is an increasing number of companies utilizing supercomputers in such diverse fields as oil exploration, VLSI circuit design, and making movies. All of these activities depend upon academic research for ideas and for trained people able to contribute effectively to the goals of the organization. In such a dynamic environment coordination is essential. This leads to the first recommendation:

Recommendation 1: NSF coordinate its programs in this area with other federal agencies and with private supporters and performers of computer-related research and education.

C. Local Facilities

Almost all scientists and engineers need a wide variety of computer services whether they use supercomputers or not. At or within close proximity of their everyday working environment, researchers need to be able to write and debug programs, store, transform, and display data, perform computations and statistical analyses, explore mathematical strategies for modelling physical phenomena, communicate with collaborators and perform many other essential professional tasks. Interactive systems are especially useful for these purposes. On them, researchers can test and make use of experimental systems. Industrial scientists and engineers have found these tools indispensable for their work. For academic science and engineering, NSF must increase its support for this purpose. The amount and type of local facilities needed will vary greatly among disciplines, types of research and working patterns of individuals. This leads to the second recommendation:

Recommendation 2: NSF greatly increase its support for local scientific computing facilities, including:

- workstations for individuals or small groups of scientists, e.g. microprocessors, graphics, data storage capability, terminals, and printers;
- computer systems for groups of researchers and graduate students;
- specialized facilities such as array processors, symbolic processors, experimental computer systems, and scientific data acquisition, storage, and dissemination centers:
- local area networks of computers and workstations.

D. Supercomputer Facilities

Since 1970, the Foundation has provided supercomputer facilities for only one community of researchers, those in the atmospheric and physical oceanographic sciences. Generally, scientists in other disciplines have been able to use supercomputers only if they performed research of relevance to the national laboratories such as Los Alamos, Livermore, NASA/AMES or Kirtland Air Force Base, or formed collaborations with foreign scientists to gain access to supercomputers in Europe.

The use of supercomputers is indispensable to the atmospheric and oceanographic scientists. For example, for the study of turbulence they need computers which are still five times more powerful than the Crays they have installed. This has led NSF to develop a plan for installing an Advanced Vector Computer at NCAR in FY 1985 which will provide that needed capability. Other disciplines have similar needs and, in some cases, problems which are even more demanding.

There is no question about the existence of problems in many fields of science and engineering which require supercomputers nor about the presence in universities of researchers interested in these problems and capable of such investigations. However, it is difficult to estimate the amount of supercomputer capacity needed because lack of opportunity has depressed this kind of research activity to artificially low levels. As a result, we cannot easily gauge either the immediate unfilled need nor the rate of growth which should be allowed if opportunity were to arise. At the Foundation-sponsored workshop on May 23, 1983, scientists in each of several disciplines described the situation in their field and estimated immediate requirements and growth rates on the basis of their knowledge of the research problems and talent available. They noted that it will take a strong sustained commitment to large-scale computing by the Foundation to persuade scientists to invest their time and effort, sometimes involving years, in preparing programs for supercomputers. Nevertheless, they could identify problems and colleagues ready to work, some of them now seeking time at foreign research laboratories. The discussions of the May 23 Workshop led to the conclusion that an immediate need exists for NSF to provide two supercomputers in addition to the Advanced Vector Computer at NCAR, and that growth by a factor of five in the following two years should be supported. This leads to the third recommendation:

Recommendation 3: NSF encourage proposals to provide supercomputer services for academic research and education in science and engineering and be prepared to support 10 supercomputer systems within the next three years.

E. Networks

1. Need for a Network

Supercomputer facilities are scarce resources. Few laboratories or universities are likely to have such state-of-the-art systems as local research tools. Rather, as in many other fields needing expensive research equipment, supercomputers will be installed as centralized facilities at a few locations. They will be open to use for meritorious research (determined by peer review) by a geographically widely scattered community of users. While some scientists may need to be physically present at the facility to carry out research projects, most will find that travel to a distant site is a highly inefficient use of their time, as well as a barrier to graduate student involvement in the research.

Remote access to supercomputers is a realistic option with existing telecommunications technology. This was emphasized by the report of the Panel on Large Scale Computing in Science and Engineering, Peter D. Lax, Chairman, [1], which made several recommendations for a national program involving advanced computer technology.

> "Underlying them all is the establishment of a system of effective computer networks that joins government, industrial, and university scientists and engineers. The technology for building networks that allow scientists to share facilities and results is already developed and understood; no time should be lost in connecting existing research groups and computing facilities."

This echoed the conclusions of the report, "Prospectus for Computational Physics", William H. Press, Chairman, [2], which recommended a computing system distributed over a network consisting of principal nodes (with supercomputer capability and large data storage facilities) linked to local nodes (with medium scale computers, advanced graphics devices, array processors, etc.). The Press report stressed the importance of networking the computer facilities and the users.

> "The difference between a computer network and a disconnected set of computers must not be underestimated. A network (and its associated software and human components) provides economies of scale in software development; it allows scientific collaborations on computational projects between participants who are physically remote from each other; it gives the user a choice among a variety of facilities that may be more suitable to his problem than the facility he happens to be physically closest to; it allows researchers at smaller universities to participate on an equal basis with larger institutions. We would not seriously contemplate the establishment of a new initiative in

computing for theoretical physics if it did not include the establishment of a supported communication network."

"There are a number of specific advantages to the distributed picture. Communication costs are, on the average, reduced by placing the physical facilities near the greatest source of their use. The network's total capacity degrades more gracefully under occasional hardware failure, and the probability of catastrophic loss of capacity is reduced almost to zero. There are often observed to be large efficiencies which occur when services such as operator time or software development are provided informally at local nodes at little or no cost (and can then filter out to other nodes). Local management and consulting services, especially when informal and on a small scale, are almost invariably more responsive to special user requirements or situations. Certain requirements (such as graphics processing, three-dimensional displays, and convenient editing on a modern screen editor) require very high baud rates. and so can only be implemented efficiently at a local site."

"The single most important advantage of a distributed network is that it can be changed gradually. It can be increased or decreased in capacity without severe dislocation of most users. It can accommodate changing technological developments without becoming 'locked-in' to a supercomputer of some one particular technological state."

2. Role of the NSF

A national computer network can be envisioned to be analogous to the interstate highway system which supplies the infrastructure necessary for transportation and communication but which leaves it up to the dynamics of the free market to work out which cities will prosper, and what their individual characteristics will be. Similarly, NSF should facilitate the use of the basic channels for computer communication between research institutions and should leave to peer review and local initiative the details of equipping individual sites, and upgrading the network capabilities available to specific nodes. Individual nodes would include large data factories where information is generated either by observation (e.g. VLA, space telescope, Fermilab...) or by computation on supercomputers (e.g. MFECC, NCAR...). Others will be smaller computational or data management facilities, institutional local area networks or clusters of work stations belonging to single research groups.

The objectives of such a network include:

1) providing access by the academic research community to multiple levels of computational resources, including supercomputers;

2) providing access to large data bases;

3) reducing wasteful duplication of effort in software development and encouraging mutual assistance and collaboration among researchers.

In initiating such a national computer network, NSF should work closely with other governmental agencies (DOE, DOD, NASA, Commerce, etc.) which share responsibility for the health of United States' science and technology. These agencies possess unique computing facilities and working networks, which should be important components in any new structure. NSF should be alert to opportunities to work with the national laboratories to provide additional remote computing capability on the network for use by university scientists, and to assist in the transfer of lab software and operating systems to interested universities.

Computer network services are an essential component of the package of actions NSF should take to improve the computing environment for research. This leads to a fourth recommendation:

Recommendation 4: NSF support networks, linking university and laboratory researchers with each other and with supercomputer centers to provide access to facilities, file transfer capability and scientific communication.

F. Advice and Oversight

Since supercomputers are a scarce resource, the allocation of computing resources should be based on need for the speed and memory capacity of supercomputers, efficiency of programs, and scientific justification of research projects. Over the years, the national laboratories have developed useful standard software packages. By sharing these software products with other supercomputer centers, the software development and operational costs of these centers would be minimized. Supercomputer centers should be operated on a cost effective basis while providing the broad range of computing services required to serve the best science. There are many potential sources for supercomputer services including national laboratories, university computing centers, commercial computer service companies, and others. There are also many possible ways for NSF to support computer services, through individual research grants, purchase of blocks of time, supporting facilities, etc. Network technology is complex and dynamic. Many new services are on the horizon and networks for research must evolve to match changing requirements with available offerings. These concerns are indicative of the large number of policy and technical quesions which NSF must decide in implementing Recommendation 1. Response to them leads to a fifth recommendation:

Recommendation 5: NSF establish an advisory committee to advise the Foundation on policies and mechanisms for supporting computer services for research and to provide oversight.

G. Academic Research in Advanced Computer Systems Design and Computational Mathematics

As shown in Section I. of this report, there are many compelling scientific problems which are beyond the reach of current computing technology. Highly parallel computer systems architectures offer a hope of addressing this need.

The advent of VLSI technology has unleashed great creative energy in academic computer research, especially in the search for new computer architectures. It is now widely believed that computer component technology is nearing the limit of speed allowed by physical principles. Further advances must come from new arrangements of the logical elements into parallel arrays of concurrently acting parts. VLSI technology makes hitherto impossible arrangements practical and offers the promise of 100 to 1000 fold increases in computer speed. Such architectures are radical departures from any we have used before, however, and require new problem strategies, new algorithms, new programming methodologies and new operating systems to make them effective.

Computer companies have shown little interest in pursuing this line of research. Most of the results have come from academia. In the last two years, however, computer and semiconductor manufacturers have formed two consortia to promote longer term research in computer systems and component technology. One, the Semiconductor Research Consortium, Inc. (SRC), is supporting academic research in semiconductor materials, manufacturing technology and circuit design. This will help make the most effective possible components available but the SRC does not intend to support research in computer systems. The other, the Microelectronics and Computer Technology Corporation, (MCC), plans to perform research in several aspects of computer technology including software productivity, packaging, CAD/CAM, and advanced computer systems design. The latter research is expected to include some work on parallel computer architectures but there is no intention of working on the design of high-performance numerical computer systems. The emphasis will be, instead, upon knowledge-based systems, data base management, and human-machine interface.

The planned DARPA strategic computer program addresses parallel architectures in part. It will support certain academic design teams working in concert with industry to build and test machines which incorporate new system concepts. It will be focused primarily on computers for artificial intelligence coupled with rapid access data handling and high speed signal processing, however, and not numerical computation which is a central concern of most scientists and engineers. As a result, it is important for the future of scientific computation that NSF play an active role in the computer design arena by expanding its support of computer research to encompass experimental testing and evaluation of design concepts for numerical computing.

On the other hand, academic researchers have a severe problem in this arena, which is to find the means to carry their computer systems design ideas beyond their desks and simulators through the stage of experimental testing and evaluation. This is an essential step but requires construction of substantial machines and the research on software systems to make them operable. Such projects are difficult in an academic environment and could benefit from industrial engineering and management experience. They are also an order of magnitude more expensive than traditional NSF academic computer research projects because they require team efforts and extensive hardware and software construction.

In closing this discussion, it is important to note that necessary as computer hardware and architecture advances are, there is an essential accompanying role and contribution of research in computational mathematics, algorithms, programming methodology and other aspects of computer science and engineering research. Indeed, past experience has shown that advances in computational capability for aerodynamics, for example, has been due in equal measure to advances in hardware and to new algorithmic approaches or problem solving strategies. Adequate attention must be given to this latter aspect of any substantial computer-related activity. Not only is it necessary that suitable relevant research activity be funded--it is also necessary that a new generation of scientists be trained in appropriate computational mathematics and computer techniques to respond to the current and future needs of this expanding technology. Consideration of all of these program elements leads to recommendation six:

Recommendation 6: Support a vigorous program of academic research and training in the areas of advanced computer systems design, including building and evaluating experimental computer systems, and in computational mathematics, software and algorithms necessary to the effective and efficient use of computer systems.

Entering the area of experimental computer systems research would present some unique problems for research support and management, however. There might be more competing academic projects than money or trained personnel available to carry them to the stage of experimental evaluation. Careful selection would have to be made to assure that only the most significant ideas, measured also against the criterion of constructibility, were selected for further work. In some cases a fusion of separately conceived academic designs might prove to be desirable. During construction, team efforts would be necessary and the use of industrial facilities and expertise would often be advantageous. Early involvement of industry would also promote technology transfer but, on the other hand, imposing proprietary controls on design information too soon would do a disservice to research progress. Finally, it would be necessary to adopt a national perspective in selecting projects for support to assure that the limited financial and personnel resources of the country are appropriately balanced among the diverse interests and priorities of DOD, NSF, DOE, NASA, NIH, industry and the universities.

IV. A Plan of Action for the NSF: Funding Implications

A. Introduction

The objective of this initiative is to improve computing services for academic research in the U.S. This divides into two subgoals: a) to improve access to modern computer technology by providing equipment, facilities, services, and networks and b) to improve computing capability through academic research into advanced computer systems and computational mathematics. NSF has a long history of support in all of these areas through research project support, equipment grants, facilities support, and national centers. Section IV. of this report provides the funding implications of a plan of action for NSF to better meet the requirements of academic science and engineering. Although this plan divides into several parts, they should be viewed as components of a package which, taken as a whole, will improve the use of computer technology for research. Each component is important and none should be asked to stand alone.

It is difficult to gauge the real needs of communities of scientists whose use of computers has been depressed by the obstacles they have had to overcome. As a result, we must depend upon comparisons with other communities which have enjoyed greater resources, insight into the important research problems and opportunities of our time, and knowledge of the research talent in the U.S. The resulting estimates are best judgments of well informed scientists and engineers, no more and no less.

One appropriate comparison is that between typical industrial or national laboratories and typical academic ones. The contrast is striking at all levels of computation services from supercomputers to personal workstations. Academic research productivity is hampered by the present level of computation support. The computation fraction of a typical NSF discipline budget is approximately 3%. It varies among disciplines, of course, being particularly high in the atmospheric sciences which support a major supercomputing facility. Not all research is as computationally intensive as the atmospheric sciences but we should remove the artificial barriers to use of computers in research in all disciplines as much as our resources will allow. It is in this spirit that the following estimates are given.

B. Budget

Note: All figures are given in FY 1984 constant dollars.

1a. Local Facilities

As noted in earlier sections of this report, researchers need a variety of computing services in their normal working environments whether they use supercomputers or not. Support required includes equipment, software, maintenance, local area networks, communication costs, technical support personnel and other costs of computer services than supercomputer time. Some of these facilities will often be used as remote workstations connected to supercomputers but most address the needs for computing which would represent an inefficient use of supercomputers but are essential to the conduct of research. NSF already responds to these needs through its research grants and equipment grants but there is a large gap between the computer support available to academic researchers and that which is considered necessary and is common in well equipped industrial laboratories. It is difficult to estimate the magnitude of this difference but it may be approached by comparing the support provided by NSF in FY 1983 with informed estimates for FY 1984, FY 1985 and FY 1986 of the actual needs of the communities served by the Foundation. The needs include workstations for individuals or small groups of scientists, computer systems for groups of researchers, specialized facilities such as array processors and local area networks. From the estimates given, budget totals for support of local facilities are summarized as follows:

Disciplinary Area	FY83	FY84	FY85	FY86
Engineering	$1.4M	$3.0M	$9.5M	$13.0M
Mathematics	1.2	2.3	6.4	7.3
Physics	2.4	4.1	13.4	10.5
Chemistry	2.1	2.7	4.2	6.4
Materials Research	2.5	3.0	5.2	6.0
Atmospheric Sciences	1.8	3.0	11.0	12.5
Astronomy	1.3	1.5	2.3	3.5
Ocean Sciences	2.5	2.8	3.6	4.5
Earth Sciences	0.8	1.0	1.3	2.5
Computer Research	13.5	16.7	25.0	28.0
Biological Sciences	1.6	2.5	4.1	4.8
Social & Economic Sciences	0.9	1.1	2.3	3.5
Information Sciences	0.2	0.3	0.7	1.2
Behavioral & Neural Sciences	1.3	1.4	1.8	3.0
Total	$33.5M	$45.4M	$90.8M	$106.7M

1b. Supercomputer Capacity

The estimates in Sections I. and III. of this report indicate an immediate need for two supercomputers in addition to NCAR's planned AVC, with growth by a factor of five in two years. As a result, the Foundation should provide two supercomputers for general scientific and engineering research in FY 84 and install eight more, including NCAR's AVC, in the FY 85/FY 86 time period.

Unfortunately, funds available in FY 84 do not allow such a rapid development of this activity. As a result, supercomputer activities in FY 84 will be limited to the purchase of services from one or more existing centers of strength. This will allow a few major research problems to be solved and program development to begin in others.

Cost estimates can be derived from the experience of the DOE, NASA and NSF national laboratories. The data in the Appendix of this report show that annual budgets for these computing centers range from approximately $6M per supercomputer to $16M per supercomputer depending upon the nature and type of other facilities and staff that are associated with the center. They also indicate some benefit from factors of scale, multicomputer centers being less expensive per unit than those which have only one supercomputer.

A supercomputer center should have a fully configured computer system--this improves efficiency as well as offering maximum capability--an adequate staff to operate the equipment efficiently and to assist users, and a complement of auxiliary facilities and services which is appropriate to the type of research computing being done. The NCAR and MFECC models fit well for single discipline centers. They each have two supercomputers with an annual budget of approximately $12M in FY 1984.

On the other hand, multi-disciplinary centers will need larger program maintenance, librarian, and user consulting staffs and, possibly, a broader range of auxiliary equipment. As a result, their costs will be higher. Also, the cost of each center will be critically dependent upon the cost of the central facility which varies from approximately $10M-$15M for currently available systems to an estimated $20M for the next generation systems.

It is clear that there are uncertainties in the actual cost of providing supercomputer services for research but reasonable estimates may be given for the near future based upon present prices. If one assumes a lease/purchase plan, the NCAR-AVC is estimated to cost approximately $6.0M/year. To this must be added staff and other operating costs which, for NCAR, totalled $5.5M in FY 1983. This brings the total to approximately $12M/year/computer for a center with one machine of the class of the NCAR-AVC. Adding a second machine would not double the cost but would increase it to approximately $20M/year for an average of $10M/year/machine. This is a reasonable average cost to use for budget estimates. On this basis and estimating NCAR's computing budget to be $12M in FY 1984 and $19M in FY 1985 if the AVC were to be lease-purchased, we find the following budget totals:

	FY 84	FY 85	FY 86
Total	\$14.4M(1)	\$70M(2)	\$110M(3)

(1) This assumes that NSF will support, in FY 84, the present NCAR facilities at \$10M and the purchase of supercomputer services from existing centers at \$4.4M.

(2) This assumes NSF will add the NCAR Advanced Vector Computer and five supercomputer facilities for academic science and engineering research to the facilities available in FY 1984.

(3) This assumes NSF will support nine supercomputer facilities for academic research and also the NCAR facilities for atmospheric and oceanographic sciences.

1c. Specialized facilities, e.g. large data banks

These are of concern to individual disciplines or subdisciplines and create an unpredictable requirement. They must be judged in terms of the research priorities at the time the need arises. Although their costs may be substantial and they may require individual accommodation in budget planning, no comprehensive estimate is possible nor is it possible to consider them adequately within this initiative. The costs of these facilities will be included in the budget plans of the disciplines requiring them as the need arises.

2. Networks

Computer networks are an essential component of an initiative to improve computer resources for academic research. They will serve three objectives:

1. provide access by the academic research community to multiple levels of computational resources, including supercomputers;
2. provide access to large data bases;
3. reduce duplication of effort in software development and encourage mutual assistance and collaboration among researchers.

Computer networking technology is already in an advanced state of development and there are many possible sources of data transmission services. NSF requested SRI International, Inc. to make a study of ways to implement computer networking for the academic research community [37]. Cost estimates can be derived from the information provided by that study. Also, DARPA, NASA, DOE and NSF have gained experience operating limited networks. In addition, NSF has discussed future network services with vendors and obtained some cost estimates from them for a hypothetical set of locations in the U.S. Combining all of these inputs, realistic investment in network hardware and software would cost:

	FY 84	FY 85	FY 86
Total	$2.1M	$7.3M	$11.5M

3. Capability - Advanced Computer Systems Research and Computational Mathematics

NSF supports computer systems research and computational mathematics through several programs. The areas of particular relevance to research computing are problem strategies or computational mathematics, algorithms, software systems and software engineering for concurrent processes, programming methodology, and computer systems design. These are supported primarily through the Computer Research Section, several programs of the Mathematics Section of MCS, the Computer Engineering program of ECSE and the Division of Information Science and Technology. Work on problem strategies is also supported through the research programs of all of the disciplines which use computers since some fraction of their investments in computer equipment, computer time, and research personnel results in new computing methodology in the respective disciplines.

A national commitment to computational support for science and engineering calls for increased academic activity in all of these areas. On the other hand, the often cited shortage of faculty and graduate students in computer science and engineering limits the resources which can be usefully invested.

In computer systems research, itself, highest priority must be given to work in experimental computer systems including the construction and evaluation of experimental hardware and software for high performance numerical computing. A program for this purpose has been described in the Lax report ([1] p. 42-76) where it is recommended that up to five construction projects be supported at any one time. Each project is estimated to cost $5M/year.

Taking all of the elements in this area of research together, a budget plan appropriate to the needs is:

	FY 84	FY 85	FY 86
Experimental Computer Systems	$1.5M	$11.5M	$23.0M
Other Related Computer Research and Computational Mathematics	$6.5M	$8.5M	$10.0M
Total	$8.0M	$20.0M	$33.0M

4. Summary

The total initiative is summarized below:

	FY 84	FY 85	FY 86
Local Facilities	$45.4M	$90.8M	$106.7M
Supercomputers	$14.4M	$70.0M	$110.0M
Networks	$2.1M	$7.3M	$11.5M
Advanced Computer Systems and Computational Mathematics	$8.0M	$20.0M	$33.0M
Total	$69.9M	$188.1M	$261.2M

C. Implications for NSF Organization

Of the several activities proposed in this initiative, the support of local facilities and of experimental computer architecture research fit well within existing NSF programs. The support of supercomputer and networking services may not. A key question in considering this issue is the nature of the supercomputer facilities. Should they be single-purpose or multidisiplinary? Should they be located on university campuses or at national centers in the manner of NCAR? There are arguments for each point of view.

1. Discipline Oriented Facilities

Discipline orientation would concentrate scientists with common scientific interests at certain shared facilities. This would encourage discourse and mutual assistance because of the commonality of experience and interests. It would also simplify the problem of peer review of applications for use of the supercomputer facilities because only problems from closely related fields would be under consideration. NCAR experience with this type of operation has been good. It is also argued that universities cannot manage large computing facilities for research well and, as a result, they should be placed in special institutes or national centers.

Another advantage of this approach is that it maps easily onto the organizational structure of NSF. On the other hand, few NSF staff have the technical competence to evaluate supercomputer or networking proposals and exercise oversight.

2. Multi-disciplinary Facilities

Some scientists, e.g. Professor Kenneth G. Wilson, Cornell University, assert quite emphatically that computer facilities organization should not reflect discipline. His argument is that in computing, attention is on the mathematical and computing methodology as well as the scientific motivation and there is much to be gained by the cross-fertilization of different intellectual backgrounds among persons having common computing interests. It is for this reason that Professor Wilson is planning an inter-disciplinary institute of theoretical research at Cornell. This commonality of methodology across widely disparate scientific applications is highlighted in a forthcoming study of the National Research Council [31] which analyzes several scientific problems from the viewpoint of the mathematics and computing algorithms involved.

Proponents of this view argue, further, that this interdisciplinary organization fits well in a university environment and, because of the urgent need to train students, facilities should be placed on university campuses where they can be accessed by students and where students can share in the cultural environment of high speed computing. They argue further, that if universities were given full support for supercomputer facilities so they would not have to justify them by their value to undergraduate education, the competing interests which have interfered with good management of computers for research in the past would not be present.

From this point of view, supercomputer and networking services are related activities, central to the research interests of all of the research directorates. They are similar in some respects to the services provided years ago by the Foundation's computer facilities program or to the synchrotron radiation facilities of today which serve chemistry, biology, and materials research.

These two examples suggest that NSF support of these activities might be centralized in one place, for example, an Office of Computing Resources, but that the specific location of that office in the organization chart of the Foundation should be of secondary importance. The office should not seem to be the captive of one interested discipline but it should be integrated with the research directorates that it serves. It should coordinate with all of the research disciplines but not interfere with already established and effective responsibility for computer activities such as AAEO's responsibility for NCAR. These criteria and the intended functions suggest creating an Office of Computer Resources reporting either to one of the research assistant directors or to the Director.

It seems probable that the impact of these questions on NSF organization is a serious concern only during a transition period from little support of large-scale computing to extensive support. If large-scale academic computing were commonplace, there would probably be a mixture of discipline oriented and inter-disciplinary facilities supported by NSF. Many NSF staff would be well versed in the culture of large-scale computing and the problems of resource allocation, proposal evaluation, and oversight could be easily shared among the research directorates. None of those hypotheses apply today, however, and the Foundation could easily impose its organizational structure on large-scale computing activity. This would be a disservice to science because we do not have reason to believe that NSF's structure applies well to computational science. For this reason, until large-scale computing grows to its natural stature in research, an office addressing large-scale (but not local) computer and networking services for all of the disciplines may be advisable.

Selected Bibliography

[1] "Report of the Panel on Large Scale Computing in Science and Engineering," sponsored by DOD and NSF in cooperation with DOE and NASA, Peter D. Lax, Chairman, December 26, 1982.

[2] "Prospectus for Computational Physics," Report by the Subcommittee on Computational Facilities for Theoretical Research to the Advisory Committee for Physics, Division of Physics, National Science Foundation, March 15, 1981, William H. Press, Chairman.

[3] "Future Trends in Condensed Matter Theory and the Role and Support of Computation," report by the Condensed Matter Theory Panel, published under the auspices of the council on Materials Science for the Division of Materials Sciences, U.S. Department of Energy, J. D. Joannopoulos and A. N. Berker, Co-Chairman, September 22, 1981.

[4] "The Numerical Aerodynamic Simulator - Description, Status, and Funding Option;" a report prepared at the request of the Committee on Commerce, Science, and Transportation, United States Senate; Dec. 1981, NASA.

[5] Numerical Aerodynamic Simulator; Background, Requirements, and Justification; NASA, Ames Research Center, Revised January 1982.

[6] "Simulation of Turbulent Shear Flows: What can we do and what have we learned?" presented by invitation at the Seventh Biennial Symposium on Turbulence, Rolla, MO, by W. C. Reynolds, Department of Mechanical Engineering, Stanford University, September 21, 1981.

[7] "A Commentary on Computational Mechanics," J. T. Oden and K. J. Bathe, Applied Mechanics Reviews 31 (1978) pp. 1053-1058.

[8] "Report on the Utilization of Large Scale Computers," Committee on the Utilization of Large Scale Computers in ASGA, Bennett Miller, Chairman, July 1, 1977.

[9] "Nuclear Weapons and Computers, Summary: The Role of Large Scale Computers in the DOE Nuclear Weapons Program," the Gold Report, prepared for the Director of Military Applications, U. S. Department of Energy, March 1978.

[10] "DOE Research in Utilization of High-Performance Computers," B.L. Buzbee, W. Jack Worlton, G.A. Michael, G. Rodrigue, LA-8609-MS, Los Alamos National Laboratory, December 1980.

[11] "A Philosophy of Supercomputing," Jack Worlton, LA-8849-MS, Los Alamos National Laboratory, June 1981.

[12] "Department of Energy FY 1983-1987 ADP Long Range Plan," Office of ADP Management, December 1981.

[13] "An Assessment of Trends in the Solid State Sciences," National Research Council, July 1980.

[14] "Long Range Plan for Nuclear Science," excerpts from a report by the Nuclear Science Advisory Committee, July 30-August 4, 1979.

[15] "Draft II - Astronomy Survey Committee: Report of the Panel on Data Processing and Computational Facilities," May 19, 1980.

[16] "Workshop on Computational Astrophysics," NASA Ames Research Center, February 1980.

[17] "Magnetic Fusion Energy and Computers: The Role of Computing in Magnetic Fusion Energy Research and Development," Office of Fusion Energy, DOE, October 1979.

[18] "Report by the Subcommittee on Computational Capabilities for Nuclear Theory to the DOE/NSF Nuclear Science Advisory Committee," December 1981.

[19] "Universities and the Future of Computer Technology Development," Kenneth G. Wilson, Laboratory of Nuclear Studies, Cornell University, Ithaca, New York.

[20] "A Data Flow Supercomputer," Jack B. Dennis, G-R Gao, Kenneth W. Todd, Computation Structures Group Memo 213, MIT Laboratory for Computer Science, Cambridge, Massachusetts, March 1982.

[21] "Foreign Travel Trip Report: Visits with Japanese Computer Manufacturers," R. Ewald, Bill Buzbee, J. Worlton, CDO/82-6782A, Computation Division, Los Alamos National Laboratory, April 5, 1982.

[22] "Supercomputers," Ronald D. Levine, Scientific American, January 1982.

[23] "An Assessment of Computational Resources Required for Ocean Circulation Modeling," Ad Hoc Committee on Computing Resources and Facilities for Ocean Circulation Modeling, Ocean Sciences Board, Commission on Physical Sciences, Mathematics and Resources, National Research Council of the National Academy of Sciences, 1982.

[24] "Nearest Neighbor Concurrent Processor," E. Brooks, G. Fox, R. Gupta, O. Martin, S. Otto, E. De Benedictis; California Institute of Technology, DOE Research and Development Report, CALT-68-867, September 1981.

[25] "Why the U.S. Government Should Support Research On And Development Of Supercomputers," Bill Buzbee, Los Alamos Report LA-UR-82-1968, July 1982.

[26] Los Alamos SCIENCE, Los Alamos National Laboratory, Vol. 2, No. 2 (1981).

[27] Justification for Class VI Scientific Computing Capability at Sandia Laboratories, Albuquerque; Sandia Laboratories, February 1979.

[28] "Trends and Pacing Items in Computational Aerodynamics," D. R. Chapman, Lecture notes in Physics, 141, Springer-Verlag, 1981.

[29] "Scientific Justification for an Advanced Vector Computer," National Center for Atmospheric Research. Prepared for NSF by the University Corporation for Atmospheric Research, December 1982.

[30] "The Influence of Computational Fluid Dynamics on Experimental Aerospace Facilities: A Fifteen Year Projection." Prepared by the Committee on Computational Aerodynamics Simulation Technology Developments, Aeronautics and Space Engineering Board, Commission on Engineering and Technical Systems, National Research Council, 1983.

[31] "Computational Modeling and Mathematics Applied to the Physical Sciences" A report prepared by the Committee on the Applications of Mathematics, Office of Mathematical Sciences, Commission on Physical Sciences, Mathematics and Resources, National Research Council, to be published.

[32] "Research Directions in Computer Engineering" Report of a Workshop, Herbert Freeman, ed., Technique et Science Informatiques, 1982, pp. 423-432.

[33] "Supercomputers: Foreign Competition and Federal Funding" Issue Brief Number IB 83102, Nancy Miller, Library of Congress, Congressional Research System, June, 1983.

[34] "Doing Physics with Computers," Special Issue of Physics Today, May, 1983.

[35] "The National Resource for Computational Chemistry, Phase 1: Evaluation and Recommendations," Report of an Ad Hoc Review Committee, Oak Ridge National Laboratory, ORNL/TM-7294, May, 1980.

[36] "The Fourth Inventory of Computers in Higher Education: An Interpretive Report," John W. Hamblen and Carolyn P. Landis, EDUCOM Series in Computing and Telecommunications in Higher Education 4, 1980.

[37] "USERNET: A Preliminary Architectural Study and Implementation Strategy," Final Report by Franklin F. Kuo et al, SRI International, Inc., June 1983.

APPENDIX

LARGE SCALE COMPUTING FACILITIES

NATIONAL SCIENCE FOUNDATION

NATIONAL CENTER FOR ATMOSPHERIC RESEARCH (NCAR)

Major Systems Currently in Place

- 2 CRAY-1A's (1 Megaword Main Memory each)
- 2 IBM-4341 Interactive Computers
- Ampex Terabit Memory Mass Store (75 x 10^9 bytes of on-line storage)

Research Community Served

- 90% Atmospheric Science, 9% Ocean Sciences, 1% other
- NCAR and Atmospheric and Related Scientific Areas in Academia
- 850 Users

Communication Systems for Remote Users

o Uninet Packet Switching Services	(1200 bps)
o Dial-up Telephone Access	(300/4800 bps)
o Dedicated Leased Telephone Lines	(2400-4800 bps)

Budget/Staff/Charges

- No charge for NSF-Supported efforts; $2200 per prime CPU hour for others
- $7.5M FY 1983 Budget-$2.0M Capital Facilities for 2nd CRAY-1A; $5.5M Operating Budget
- Staff of approximately 100

Plans

- FY 1984-New Mass Store System ($3M)
- FY 1985-Class VII Computer System (also to serve ocean scientists) ($20M)
- FY 1986-Wide-Band communications ($17M)

NATIONAL MAGNETIC FUSION ENERGY COMPUTER CENTER

Major Systems

- CRAY-1 (1 Megaword Main Memory); CRAY-1S (2 Megaword Main Memory)
- CDC 7600
- CDC 38500 Mass Store (62,500 Megabytes of storage, retrieval time is 10 sec.)

Research Community Served

- 95% to Magnetic Fusion Energy Research Community
- 5% to Other Energy Research
- 2,000 Users

Communication System for Remote Users

- Dual 56K bps Satellite System with 4 remote User Service Centers
- Dial-up Service through TYMNET, ARPANET, Commercial, or FTS
- Dedicated Leased Telephone Lines (4800 bps)

Budget/Staff/Charges

- No Charge
- $12.5M FY 1983 Operating Budget (Includes lease of 1 CRAY)
- Staff size approximately 85

Plans

- Class VII Computer System in FY 1984
- Another Class VII System in FY 1987
- Modest Expansion of Mass Store in FY 1984

LOS ALAMOS NATIONAL LABORATORY

Major Systems

- Access Partitioned into Three Levels: Administrative, Secure, Open
  - Open: 1 CRAY, 3 CDC 7600's 1 CYBER 825
  - Total: 5 CRAY's, 4 CDC 7600's, 3 CYBER 825's
- IBM 3850 Mass Store Cartridge System (1.2 trillion bytes, 85% returned in 5 - 10 secs.)

Research Community Served

- Government Agencies, Laboratories, and Non-Profit Institutions
- 4400 total Users

Communication System for Remote Users

- Dial-up Telephone Access through Arpanet, Telenet and Commercial Lines (300/1200/2400 bps)
- Dedicated Leased Telephone Lines (Up to 5600 bps)
- Magnetic Fusion Energy Network (56K bps)

Budget/Staff/Charges

- $39M Budget for FY 1983
- Staff of 240 FTE's
- $636 charge for each prime time CRAY CPU hour. Also there are charges for Connect Time, Ports, Storage, Output.

Plans

- Install Denelcor HEP-1 Experimental System in FY 1983
- Upgrade a CRAY-1 to a CRAY-XMP in early FY 1984
- Add 3 Class VII Systems (FY 1985-1988)
- Upgrade Mass Store System in FY 1984, FY86, FY88

NASA - AMES RESEARCH CENTER

Major Systems

- CRAY - 1S (1 Megaword Main Memory; 16 Megaword Solid State Disk Memory comparable in response time to Main Memory)
- CDC - 7600

Research Community Served

- Computational Fluid and Aerodynamic Community*
- In-House Computational Chemistry, Astrophyssics, Stratospheric Modeling and Image Processing Communication for Remoter Users
- 500 Users

Communication System for Remote Users

- Dial-up Commercia and FTS Telephone Lines (300/1200 bps)
- Arpanet (3-6 Kbps)

Budget/Staff/Charges

- FY 1983 Budget at $9.0M with a staff of 100+
- No direct charges to NASA grantees
- NASA donates time for cost sharing purposes on projects of joint interest with other agencies
- $2,000/normal priority CPU hour for CRAY; $750/normal priority CPU hour for 7600

Plans

- Mass Store System in FY 85; Upgraded to maximum configuration in FY 87
- Upgrade CRAY-1S to CRAY-XMP in FY 84
- FY 87 goal of serving the Computational Fluid & Aerodynamics Community with:
  - 2 Class VII Computers
  - 1 Class VI Computer (enhanced)
  - Remote Users Acessing Systems through Satellite Based Sending and Receiving Communications

*In conjunction with joint research activities conducted at NASA-Ames.

NASA - GODDARD

Major Systems

- CYBER 205 (2 Pipelines, 1 Megaword Main Memory)
- IBM 3081
- Amdahl V6, V7B
- Masstor 861 Cartridge System (100 billion bytes; access time a few seconds)

Research Community Served

- NASA-Goddard Scientists
- NASA Funded Researchers Working in Earth Science Observations
- Researchers working with NASA on common areas of interest
- 1000 users

Communications for Remote Users

- Dial-up telephone lines (300 - 4800 bps)
- Dedicated leased telephone lines (up to 4800 bps)

Budget/Staff/Charges

- $10M FY 1983 Budget (Capital and Operating Costs)
- Staff of approximately 60 FTE's
- No direct charges to NASA grantees working on Earth Observations
- To others, the charges are approximately $1000/CPU hour for the CYBER 205, and $500/CPU hour for the IBM 3081. In the future there may be charges for Connect Time, Storage and Output.

Plans

- Node on NASA Wide-band Communications System Linking Large NASA Computers
- Massively Parallel Processor (MPP)
  - 128 x 128 Array Processor; 10 Million Instructions/Sec.
  - Built under Contract with Goodyear Aerospace
  - Cost under $7M
  - To be delivered in late FY 1983

NASA - LANGLEY

Major Systems

- CYBER 203 - (1 Megaword)
- CYBER 730, CYBER 170's, CYBER 750
- Masstor 860 Cartridge System (80-100 billion bytes of storage)

Research Community Served

- NASA-Langley Research Scientists
- NASA Funded Researchers working in structures, strength of material, aeronautics and space, and atmospheric modeling
- 250 Users on CYBER 203; 2000 overall

Communications for Remote Users

- Dial-up telephone lines (1200) bps

Budget/Staff/Charges

- $16M for FY 1983
- Staff of approximately 350
- Charge $1300 per CYBER 203 CPU hour

Plans

- Upgrade Memory on CYBER 203 to 32 Megaword
- 50 Megabit Local Communication Network Langley

NASA-LEWIS

Cleveland, Ohio

Contact: Dr. Ford, Director for Scientific Computing

Major Systems Currently in Place

- CRAY 1S; in place 9 months; used at 50% capacity
- IBM 3033--370 (attached processor)
- UNIVAC 1100/42

Research Community Served

- Theoretical and experimental scientists and engineers
- Mostly aerodynamics and related rields
- Available to NASA LEWIS scientiists and engineers, other NASA Centers, NASA contractors, other government agencies, and a few universities

Communication System for Remote Users

- Dial-up telephone access
- Dedicated leased telephone lines
- CAT TV cable--main access, very broadband cable, 8 megahertz, allows great variety of networks

Budget/Staff/Charges

- Free for NASA-supported efforts
- No formal mechanism for charging "other" users; there will be one in place shortly; ballpark figure of $1,000/hour
- 2,500 people employed at computer center of which approximately 1,500 are regular users
- FY 1983 computer center budget-larger than NCAR

Plans

- Acquire two new scientific computers and one new computer for business and information management; IBM 3081's
- Plans to upgrade memory
- No plans to acquire another supercomputer

COLORADO STATE UNIVERSITY

- Leasing a CYBER 205 ($2.3M/Year)
  - o 2 megaword central memory
  - o 2 vector pipelines
  - o CDC 819 mass store system (3 billion bytes of on line storage)
  - o Cost for university users is $2000 per hour

- CDC Contributions
  - o Purchased $1.7M worth of CPU time
  - o Distributed time to 30 individual scientists from 22 universities for independent scientific research

- Examples of Research Projects using CYBER 205 are:
  - o Petroleum Engineering & Reservoir Analysis Problems
  - o Subsurface Geological Mapping
  - o DNA Modeling
  - o Monte Carlo Techniques Applied to Plasma Physics

- Access
  - o Dial-up telephone lines
  - o CSU Network - dedicated lines

U. of Mass	14.4 Kbps
Dartmouth T.S.	9.6 Kbps
National Bureau of Standards	9.6 Kbps
Solar Energy Research Institute	9.6 Kbps
Several Petroleum Research Companies and Centers	9.6 Kbps

  - o National Bureau of Standards

    Working agreement for NSB researchers in mathematics and computer science

  - o 20% of CPU utilized

UNIVERSITY OF MINNESOTA

- Leasing CRAY 1A, $1.56M/year
  - o 1M words of memory (maximum capacity)
  - o 3 disc drives
  - o Used at one-third of real capacity
  - o Front-end disc capacity of over 12 gigabytes
  - o Front-end interactive computers: VAX 11/780 and CYBER 730

- CRAY Contributions
  - o System purchased at a $2M discount
  - o 5 years maintenance at nominal changes
  - o CRAY Inc. is allotted the night shift (7 hrs.) for development work

- Cost of Services:
  - o Internal rate--at $2,000/hr.--available to university personnel, federal, and state government and other non-profit organizations
  - o External rate--mainly to local high-tech industries--50% mark-up ($3,000/hr.)
  - o Available to EDUNET network
  - o The above rates subject to substantial discounts (as high as 50%) for high-volume users

- Examples of Research Projects
  - o Chemical engineering--phase changes in material
  - o Petroleum recovery
  - o Finite element computation
  - o Image processing--animated movies involving the studies of molecules
  - o Economic modeling

- Future Plans
  - o Replace front-end
  - o Proposal to form a consortium with four industrial users for a one-fifth share in a CRAY XMP untial a CRAY II is available.

- Access
  - o Dial-up telephone lines
  - o Dedicated lines 1200 Bps, 4800 Bps, 9600 Bps, 19.2 Kbps

PURDUE UNIVERSITY

- Acquired a CYBER 205 for $5.5M
  - o Funding from University and State of Indiana
  - o 1 megaword central memory, 2 vector pipelines
  - o At the moment, there are a limited number of production runs, much work is occurring on software and systems development, approximately 6 CPU hours per day are available to users

- CDC Contributions
  - o Purchased $1.5M worth of CPU time for computer-aided design research

- Plan for Enhancement
  - o Increase memory to 2M words ($2.3M)
  - o Add 2 high speed channels, disk and tape systems, graphics systems ($1.8M)

- Some examples of research projects
  - o Biological molecular structures
  - o Solar cells
  - o Mesoscale storm models
  - o Chemical molecular dynamics
  - o Solution of certain elliptic PDE's

- Establishing a Center for Vector and Parallel Computing
  - o Research in algorithms and software for Vector and Parallel Computing

BOEING COMPUTER SERVICES

- Owns CRAY 1S
- Used at 50% capacity
- Direct dial-up lines
- Dedicated lines

CYBERNET

- Acquired CYBER 205
  - o time sharing option
  - o cost depends on type of work load
  - o types of applications include weather predictions and nuclear codes
  - o dial-up telephone lines
  - o use at one-half of full capacity

UNITED INFORMATION SERVICES

- Own CRAY 1S
- Time sharing and batch
- Applications include high-tech companies, business, i.e. insurance groups
- Cost depends on I/O units used and type of peripheral equipment
- Dial-up lines
- Dedicated lines 1200 bps - 9600 bps

CORNELL

Equipment

IBM 370-168, several 4341's (central campus facility), IBM 3081 being installed
FPS AP-190L Array Processor
FPS-164 Attached Processor

No. of users (Array and Attached Processors) 50-100 direct, available to 2000 general users

Costs

Array Processors: capital costs: $0.5M
Array Processor operating expenses and support costs: $150K/yr, staff of 3+ (120 in Computing Center)

Array Processor typical rates: $20-$40/CPU hr non-profit
$75/hr external

Scientific computing charges on IBM hosts: $500-1500/hr on IBM 370/168
Computing center budget $10M

Access available through TELENET, TYMNET, BITNET.

Aims

Cornell measures the success of its scientific computing facility by the quality of research that individual faculty members achieve using the facility. At present a consortium of four faculty research groups are the primary focus of the scientific support effort:

Keith Gubbins (chemical engineering, molecular dynamics)

Stu Shapiro and Saul Teukolsky (astrophysics and general relativity, fluid dynamics, particle dynamics)

Kenneth Wilson (high energy physics, critical phenomena, Monte Carlo simulations)

G. Chester, N. Ashcroft, J. Wilkins, (solid state physics):

The computing capability available to each one of these four research groups is approximately 2,000 hours/year of processor time, which is competitive with the capabilities per research group supplied at major national facilities.

Future Plans

Obtain a 4000 Array Processor facility, organized with high speed parallel interconnects between processors, to achieve 30-100 Gigaflop* capability. Target date: Fall 1986

One or two 64 processor facilities to be acquired earlier for benchmarking.

A joint Cornell-computing industry proposal to realize these plans should be ready by Fall 1983.

This facility will provide a desperately needed thousand-fold increase in capability for leading-edge research projects at Cornell.

In addition a distributed supercomputing network of superminicomputers, attached processors, and other high performance systems is being established to expand Cornell's capacity at present levels of capability. The aim is to serve the entire theoretical science, theoretical engineering and theoretical agricultural communities at Cornell, totalling some 500 individuals.

Also, use will be made of national facilities, starting with the Goodyear MPP at NASA Goddard Laboratory.

Emphasis will continue to be on development of new, cost-effective approaches to Scientific Computing.

---

*1 Gigaflop = 10^9 arithemetic operations/second

1884 1984
A CENTURY OF ELECTRICAL PROGRESS
IEEE

The United States Activities Board

SCIENTIFIC SUPERCOMPUTER COMMITTEE
REPORT

Edition of
October 25, 1983

Chairman: Dr. Sidney Fernbach

CONTACT: Edith T. Carper
IEEE Washington Office
Suite 608
1111 19th Street, N.W.
Washington, D.C. 20036

(202) 785-0017

SUPER SCIENTIFIC COMPUTER COMMITTEE MEMBERS

Dr. Sidney Fernbach, Chairman
Consultant

Dr. Alfred E. Brenner, Fermi National Accelerator Laboratory
Dr. James C. Browne, University of Texas
Mr. Harvey G. Cragon, Texas Instruments
Dr. James F. Decker, U.S. Dept. of Energy
Dr. Tse-yun Feng, Ohio State University
Dr. Frank Kuo, SRI International
Dr. Lawrence Lee, National Science Foundation
Maj. Carl Edward Oliver, U.S. Dept. of Energy
Dr. John P. Riganati, National Bureau of Standards
Dr. Paul Schneck, Office of Naval Research
Dr. Daniel Slotnick, University of Illinois
Mr. Hugh F. Walsh, IBM

ADVISORY MEMBERS

Mr. Theodore H. Bonn, Sperry Defense Electronic Systems Operation
Member, IEEE Board of Directors

Dr. Jack Lubowsky, IEEE Congressional Fellow
Senate Governmental Affairs Committee
Subcommittee on Energy, Nuclear
Proliferation and Government Processes

Mr. Doyce Satterfield, IEEE Congressional Fellow
Sen. Howell Heflin's Office

IEEE STAFF

Ms. Edith T. Carper, Executive Secretary
Ms. Sally A. Gustafson, Public Affairs Secretary

EXECUTIVE SUMMARY AND RECOMMENDATIONS

Supercomputer application and development in the U.S. have reached a crossroads. The U.S. must take giant steps forward or relinquish leadership. We must improve supercomputer performance by recognizing that our rate of improvement has slowed down. We must persuade U.S. policymakers that advanced supercomputer technologies are important in solving pressing technology problems that range from automotive and aircraft design to the most difficult problems of high energy physics.

We must respond to foreign challenges as well. Active development programs are in progress in Japan, Great Britain, France and other countries but the principal challenge comes from Japan. Evidence of the foreign challenge was the 1981 announcement in Japan of concurrent commercial efforts to build computers to outperform the Cray 1 and the CDC Cyber 205, along with Japan's National Superspeed Computer Program to develop new components and architectures for the 1990's.*

The United States Activities Board of the IEEE, with an awareness of its responsibilities as a major technical/professional society, organized a committee to make recommendations for correcting the U.S. course. The Committee, chaired by Dr. Sidney Fernbach, consists of seventeen individuals from industry, academia and government who represent all aspects of the supercomputer community (see biographies in Appendix II).

The Committee found that the current market for supercomputers is small but expanding, and there is increasing awareness of the potential for their use. It views the emerging market place as extremely important and potentially very large.

In addition, the economic and national security importance of supercomputers is far in excess of the number in use at this time. Each generation of mainline computer systems has descended from the technology of the previous generation of supercomputers. Without supercomputer leadership, U.S. computer vendors will soon be at a competitive disadvantage. Further, the frontiers of science and technology often yield only to the power of a supercomputer. Finally, the applications of supercomputers to weapons design and to military intelligence alone are critical to the U.S. security.

If the nation is to protect its military and economic security, and if industry is to maintain its technological vitality and leadership, lower-cost and significantly better performing supercomputers are required in this decade as well as on into the 21st century.

* The most relevant programs in Japan are the three supercomputer development programs under way at Fujitsu, NEC and Hitachi, all with design objectives of at least the levels of the Cray 1/Cyber 205; the MITI architecture and technology support program for subsequent (1990s) supercomputers; and the distinct, longer term, Artificial Intelligence-based Fifth Generation Computer Program described in Appendix I.

For industry, supercomputing has become an enabling technology that is replacing traditional "experimental" methods in engineering with computer simulation. This transformation leads to better design optimizations, faster development cycles, and lower development costs. The opportunity to use this technology to increase productivity and to sharpen the competitive edge of U.S. industry is profound.

The most advanced machines becoming available today have taken a step in a new direction by incorporating several processors into one system. Future generations of supercomputers will continue this evolution into systems of many processors, requiring new developments in software to make them more effective. Early user familiarity with these new approaches is essential to maintain the U.S. lead.

Many aspects of technology which have the potential for large increases in scientific supercomputer capability are not being developed in the U.S. as rapidly as they are in other countries. These include proposals and ideas for performance increases through the use of new semiconductor components, new packaging technologies, and high performance peripherals that are not being pursued here with the vigor required. Much U.S. supercomputer development is now dependent upon Japanese components; and, because supercomputer manufacturers are not vertically integrated in the U.S., as Japanese manufacturers are, both advanced components and peripherals for supercomputers are more difficult to obtain. To achieve higher performance levels, special attention and support must be provided to all elements of supercomputer technology.

We believe that approximately $100 million a year for five years will help provide for the deficiencies of the past and bring supercomputer development to a reasonable state-of-the-art level at the end of this period of time.

Only a Federal Government program can have the scope and resources to meet the challenges of foreign governments and industries determined to wrest leadership of these critical technologies from the United States.

The IEEE United States Activities Board recommends that the Federal Government,..

1. make a commitment to maintain leadership in scientific supercomputer development with a well defined, long range national program.

Advances in computer technology that have spawned the development of today's scientific supercomputers have been critically important to the United States. The supercomputer has been the major factor in maintaining U.S. leadership in other technologies and vital scientific areas. Without question, tommorrow's supercomputer is necessary for continued significant advance in science and technology, increase in economic productivity and improvement in international trade.

2. encourage both the continued development of scientific supercomputers and their availability to industry through the research community by adopting such measures as

   * direct support for development of new supercomputer generations,

   * direct support for establishing and operating several supercomputer centers for research, teaching and applications development,

   * tax relief or other financial incentives to supercomputer and related manufacturers,

   * tax relief or other financial incentives to supercomputer buyers and users,

   * increased direct support of related high risk research,

   * anti-trust relief to the functioning of industrial R&D consortia.

Since one-half of today's supercomputer systems in the U.S. are located in laboratories under the supervision and control of the Federal Government, it is likely that the Government will be a major consumer of future supercomputer systems. This limited marketplace may not offer necessary incentive for individual firms to invest aggressively enough to assure our leadership in future generation computer systems. Therefore, the Government should purchase (through its agencies) a minimum of 12 * highly advanced scientific supercomputer systems specified to have significantly increased performance. This initial direct Government purchase would contribute to establishing a sufficient market to manufacturers to produce such systems.

The Government should provide tax incentives that (1) stimulate companies to become and/or continue to be developers and producers of scientific supercomputer systems, (2) stimulate companies to become and/or continue to be developers and producers of associated software, component technologies, peripheral equipment, and algorithm/computational mathematics, and (3) stimulate users of supercomputer systems to acquire the systems. Such tax relief measures would allow architectures and component technologies to be developed that would match the applications of existing and needed scientific research programs.

The Government should provide direct funding to industry and university laboratories with advanced research programs in hardware and software, to ensure continuing technology innovations necessary for supercomputers of the future. Also, the Government should exempt industrial and/or university consortia from existing strictures of the Government, such as anti-trust, that inhibit the combining of resources to address these problems: too large for any one company to attack.

* To correspond with the number of major government research laboratories now using supercomputers.

3. support of all technologies needed for the development of supercomputer systems.

The major technologies required are:

- components
- architecture
- system software and applications,
- and peripherals

4. designate a lead organization to coordinate the roles of the many Federal agencies which are critically dependent on supercomputer systems.

Just as NASA was given the responsibility to lead in the national commitment to be foremost in space, responsibility for scientific supercomputer development leadership should be assigned to an organization with significant experience and a strong interest in continuous and rapid evolution of scientific supercomputers. The criteria for selecting this organization should include supercomputer experience with respect to: use, management, development of software and systems technology.

The organization ideally ought to exhibit a broad range of supercomputer needs and commitment to share supercomputer experience and technology in an appropriate manner.

The existing FCCSET Subcommittee on supercomputers of the OSTP could serve as an oversight committee to mandate responsibilities to other organizations/agencies. Since the agencies with strongest expertise in, and utilization of, these numerical machines are DOE, NASA, and NSF, one of these agencies should be selected to lead the national program. DOE has a long and successful history of bringing new supercomputers on line and into productive use.

5. establish and extend national performance goals for the development of scientific supercomputer systems well into the 21st century. These long term goals should reflect the needs of Federally sponsored scientific research, both current and projected.

The Federal Government, through its many agencies, is investing heavily in microelectronics, software, and general computer technology. However, there is no national policy that outlines computational goals and how these many programs are to be coordinated to achieve these goals. Unquestionably, many scientific programs have needs that far exceed the capability of today's supercomputer systems and, thus, a coordinated national program specifying goals and objectives is urgently needed. These national goals should also reflect the needs of non-federal organizations conducting scientific research.

Maintaining U.S. leadership in the development and use of scientific supercomputer systems will, because of international competition, require a well planned and long-range commitment by the Government. This long-range program must ensure that both proper goals and incentives are available to stimulate the development of all elements of the scientific supercomputer system.

INTRODUCTION

The United States is continuing a gradual transition from an economy based on industrial production to one based increasingly on high technology and information services. Other industrial nations are involved in the same process. In 1981, Japan's Ministry of International Trade and Industry (MITI) announced the formation of two major computer projects: one on supercomputers and the other on the development of fifth-generation computer systems. These projects represent serious challenges to the U.S. superiority in computers and information processing. If left unchallenged, Japan's major initiatives in computer technology could result in the U.S. losing its number one position in computers to Japan. European countries have also recognized the challenge and have adopted plans to meet Japan's initiative.

Maintaining our position in computer technology is fundamental to the revitalization of the entire industrial economy in the U.S. One of the most important computer technology areas is that of supercomputers. These are computers which are the fastest and most powerful computing systems available at any given time. The development of supercomputers significantly influences the development of less powerful computers. In addition to the Japanese National Superspeed Computer Program to stimulate component and architectural developments for the 1990's, three major Japanese companies are currently manufacturing supercomputers at or beyond the level of current U.S. systems.

Reports from the Lawrence Livermore National Laboratory and the Los Alamos National Laboratory concerning a visit to Japan's computer manufacturers (in early 1982) led to a study under the sponsorship of the National Science Foundation and the Department of Defense with the participation of DOE and NASA. The result was the Lax report* which appeared in December, 1982. It made three recommendations, chief among which was one urging the Federal government to provide much greater supercomputer access to the academic research community.

This was followed by a study requested by the Office of Science and Technology Policy of its FCCSET (Federal Coordinating Committee on Science, Engineering and Technology) committee, to provide guidelines for promoting 1) supercomputer development, 2) supercomputer access, and 3) research coordination. Several of these studies have been completed but are not yet available.

In addition, NSF convened an internal committee to formulate budget plans to implement the Lax report's proposal for supercomputer access for the academic research community. The committee's report is currently being reviewed by OMB.

---

* Report of the Panel On Large Scale Computing in Science and Engineering, Peter D. Lax, Chairman (Dec. 26, 1982)

On the Congressional level, hearings on supercomputers have also been held. In addition, the Office of Technology Assessment (OTA) held a workshop, in July, to study the problem. Its report should soon be issued.

The IEEE United States Activities Board strongly believes that U.S. policy makers in government and industry must establish a national policy for meeting the foreign challenge. In this, the IEEE report, which is aimed at senior U.S. industry and government policy makers, the authors present the key underlying issues behind the economic and technological challenge posed by competitors in the supercomputer arena. The report delineates the U.S. needs for scientific supercomputer systems, defines the technological requirements for the development of future supercomputers, and examines world competition in this arena. The report concludes with in a set of policy recommendations for the Federal government which, if accepted, would result in a coordinated plan of action for the U.S.

To help in understanding the supercomputer situation in the United States, we give the following short history. For many years Control Data Corporation was the chief vendor of supercomputers, starting in 1964 with the 6600. CDC in 1969 produced an improved version, the 7600. In 1974 it introduced a vector machine, the STAR 100 which was later (1981) upgraded to become the Cyber 205. Approximately fifteen Cyber 205's have been delivered in the past few years.

In the late 1960's IBM mounted some competition for the 6600 and 7600 but then dropped out of the supercomputer arena. Texas Instruments built seven ASC supercomputers (Advanced Scientific Supercomputers) but also dropped out of contention. Burroughs built one ILLIAC IV, then designed a BSP (Burroughs Scientific Processor), which was built and scrapped. The only real competitor for CDC ultimately was Cray Research, a company organized in 1972 by Seymour Cray (who had designed the 6600 and 7600 for CDC and then left to start his own company). Its first computer, the Cray-1, was delivered in 1976. Currently, approximately 60 have been installed. Both Cray Research and CDC are now designing more advanced systems intended for delivery by 1986.

A newcomer to the scene is Denelcor. Although its first attempt, a multiprocessor, the HEP-1, may not be considered a supercomputer by some (because of its relatively low performance), its new product should be competitive with those of Cray Research and CDC. More details on these systems appear in a later section, on The Competitive Threat.

APPLICATIONS OF SUPERCOMPUTERS IN INDUSTRY AND GOVERNMENT

Introduction

Supercomputers are the most advanced numerical computers in the world. They are especially well-suited to handle computationally intensive problems - ones that involve numerous variables and require billions upon billions of calculations for solutions. While Japan and Western European countries have made substantial progress in using supercomputers, this nation continues to be the international leader in their application. Evidence of this is profuse and can be categorized into five areas. These are:

- National Security. Tremendous progress has been made in the design and computer simulation of weapons systems and in intelligence data analysis and cryptography.

- Energy Development. This nation's fusion energy programs have advanced rapidly by using supercomputers to simulate behavior of thermonuclear plasmas and to design future experiments. Industry has successfully used supercomputers to assist in discovering petroleum and mineral deposits and to maximize petroleum recovery.

- Health and Safety. This includes the development of safety systems to prevent nuclear reactor accidents, the creation of organisms to contain and eliminate hazardous wastes, and the detection and forecasting of severe weather.

- Scientific Research. Advances in the studies of space science, atmospheric and oceanic phenomena, and theoretical physics can be directly attributable to the availability of supercomputers.

Thus, supercomputers have become critical in solving the complex problems vital to our security, economic and scientific leadership, and the health and safety of our citizens. Even so, the full potential of supercomputers is just starting to surface. In the 1990s many new important and difficult problems will be amenable to solution, but only if vastly more capable supercomputers are available. A discussion of some current applications in industry and government follows.

Supercomputers in Industry

Starting in the late 1970's, applications requiring and justifying a supercomputer have developed and spread into the industrial sector and the number of users whose work demands this level of performance had begun to increase dramatically. In many instances this expansion has been related to the change in the value or cost of energy.

Examples are manifold. In the petroleum industry, the value of a new oil field discovery has caused extensive activity in exploration giving a dramatic rise in the amount of seismic analysis being performed - an application well suited to the characteristics of available supercomputers. The cost of an entire supercomputer installation can be recovered by avoiding the drilling of even a single dry well. Similarly, the value of extracting an additional few percent of oil from existing reservoirs easily covers the cost of a supercomputer by one or more orders of magnitude. Again, this application - reservoir modeling - is well suited to solution on supercomputers.

In such industries as automotive and airframe manufacture, the need for increased fuel economy has led to extensive redesign. This has led to a tremendous increase in the amount of computation related to structural analysis and fluid dynamics; structural analysis caused by the need to redesign smaller, lighter and yet safer vehicles and airframes using newer, lighter materials and fluid dynamics because of the need to reduce the losses due to air resistance and turbulence. The economic and fuel cost advantage of fuel efficiency improvements as low as two or three percent are significant. Computation on supercomputers is also required for the design of lighter, more efficient engines, both piston and fan-jet.

The nuclear power industry is another aspect of this nation's attempt to increase the supply and efficiency of energy resources and thus reduce the cost of energy. because of the risks and costs involved in experimental nuclear reactor systems, this industry has turned to computer modeling of reactor designs and designs of safety and emergency systems. It is also using linear programming techniques on supercomputers to optimize power grids.

The relatively recent availability of low-cost, high density electronic circuitry has caused the designers of a wide variety of products to turn to these devices for better control and added function. The design of these L.S.I. (and soon V.L.S.I.) circuit chips is not confined to the electronics industry even though their manufacture is. Thus a number of industrial organizations are using large amounts of their own computing capacity or buying time on supercomputers offered by service bureaus to design custom-made chips.

The large scale computer manufacturing industry is turning to parallel processors to develop the computational power needed in future supercomputers. As a result of physical constraints, the single processor can have only limited performance improvements. Today's supercomputers are being used to emulate models of proposed parallel architectures and to develop and evaluate algorithms, systems software, and interfaces to other system components. More powerful supercomputers are required for designing the higher performing systems needed in the 1990s and beyond.

The area of computer aided design of machine parts and the testing of stress and strain properties of structures are examples of important applications within the U.S. manufacturing industry. The design process consumes much of a product's development costs. Design automation, including placement, layout, simulation, testing, and performance verification is crucial to the economic health of the American manufacturing industry.

The manufacturing industry is also benefiting from the rise of supercomputers in robotics research. Specific research goals are directed toward visual sensing systems, high speed image processing and recognition systems, robot control systems, and three dimensional object modeling systems.

Industry is using supercomputers for generating state-of-the-art computer graphics imagery and animation. The goal is to develop a new level of realism in special effects feature films. Depending on subtleties of motion, expressions, coloring, etc., it may take as much as 500 billion arithmetic instructions to produce one second's worth of motion graphics. More powerful supercomputers will increase quality at an overall lower cost of production.

While the chemical and drug companies do not own supercomputers, the larger companies within the industry are considering their acquisition. If they were acquired, this industry would use them for research on problems in chemical separation and purification, molecular dynamics, and for modeling of drug action and interaction. These applications are constrained by the unavailability of ultra high performance, reasonably priced computers.

Overall, the current population of supercomputers in industry is limited by two major factors. The first is cost and the second is application availability. The two feed on each other. Reduced cost attracts application support, which expands the market and further reduces costs. The availability of application packages at reasonable cost has been a significant factor in the great increase in the general purpose computer market in the last 15 years. The same effect can occur in the supercomputer market.

Supercomputers in the Federal Government and the University

While the industrial sector has recently begun to use supercomputers for research and design, the Federal Government continues to use them for defense, space, energy, and science research projects.

The nation's weapons program depends in part on computer modeling of nuclear testing. It has been apparent for many years that the ability to produce smaller, safer, and more efficient devices is critically dependent on creating and executing better computational models. Furthermore, by using supercomputers for modeling nuclear weapons performance, weapons designers are able to conduct numerous computer experiments, thus reducing the need for the cost and complications of field testing.

Researchers at national laboratories and universities use supercomputers for the study of advanced energy technology. Much of this research has been in the behavior of plasma as it interacts with magnetic fields and of fusion fuel when compressed under intense pulses of energy. This research could ultimately lead to the development of fusion power reactors. Supercomputers have permitted scientists to simulate many high energy physics reactions more accurately, but new and improved models are needed.

As aircraft and aerospace vehicles grow more complex, the need for more sophisticated and extensive design tools has increased correspondingly. The cost of performing wind-tunnel and experimental flight tests has risen significantly as the costs of labor and energy escalate and as aircraft designs have become increasingly complex. Supercomputers are used in computational aerodynamics to simulate aerodynamic flow fields through the numerical solution of fluid dynamic equations. The potential benefits derived from improved computational aerodynamics lie in all phases of aeronautical research and development from basic fluid dynamic research to aircraft configuration refinement, optimization, and performance verification.

During the past two decades, the U.S. space science program, through the development and deployment of remote sensors, has been gathering enormous data sets on the earth and solar system. Advances in sensor technology, communications, and computational capabilities have been important ingredients in this program. Within the space science data management environment, supercomputers are one of the essential elements. The steady increase in data volume, variety of space-derived data, subsequent processing requirements, and modeling needs of scientific investigations have greatly increased the need for the fastest and most capable computers.

The dynamic nature of the atmosphere and its effect on human activity has led to the development of complex technologies in weather observing and forecasting capabilities. The supercomputer is an important part of this technology. The National Weather Service depends on today's supercomputers for the processing and dissemination of meteorological data and for the developing of forecasts. In addition research in atmospheric and oceanic sciences involves analysis of very large quantities of data and the development and analysis of increasingly sophisticated deterministic models of the atmosphere and oceans. Supercomputers play a central role in this research.

In recent years scientific research at universities has greatly increased in the areas of numerical simulation of physical events and of longer term variations in nature. Examples of these are turbulence, global atmospheric and oceanic circulation, structure of liquids and their surfaces, heavy-ion collisions, etc. Much of this activity is a result of availability of low cost "mini" computer systems. Nevertheless, many university researchers have found that in order to simulate these events more realistically, they must turn to the use of supercomputers. Examples abound in the fields of materials research, structural mechanics, chemical engineering, atmospheric sciences, physical oceanography, and theoretical physics, including lattice gauge theory. Prof. Kenneth Wilson received the 1982 Nobel Prize for his computational work in lattice gauge theory.

There is a need for federally supported researchers to develop more complex and realistic models. This can be accomplished by considering models with greater dimensionality, i.e., three dimensional space plus time, higher grid point and time step resolution, and more complete representation of the laws of physics. Thus to solve the problems that defense, space, energy, and science programs face in this decade, a very substantial increase in computational speed and memory size of supercomputers is needed. This is described more completely below.

The Requirements for More Powerful Supercomputers

Supercomputers are most frequently used for simulating complex events which occur in nature. These natural events are often so complex that it is difficult to develop sufficiently accurate simulations or "models" of them. There are two main reasons for this: (1) as the understanding of physical processes improves, the models used to emulate them necessarily become more complex; and (2) the currently available large-scale scientific computers such as the CYBER 205 and the CRAY X-MP are not fast enough and do not have sufficient storage for production runs of complex models. Even to solve the problems they face in the decade of the 1980s, defense, space, energy, and science programs need an increase of about a factor of 200 in the speed of computation in this decade if the nation is to protect its military and economic security..."and to do the research necessary to maintain the scientific leadership that is key to our security in the future." Similar increases are needed in the 1990s.

In order to understand the requirements for more powerful computers, we must explore the generic reasons for having increased computational speed and storage. In this context, the term "complexity" will be used to refer to all causes of increasing requirements for computational speed and storage.

Generic requirements for faster computers

Dimensionality. The real world exists in three space dimensions plus time. If computational models reflected the real world exactly and completely, they would treat all four of these dimensions and other characteristics such as energy and fixed angles that are equivalent to additional dimensions. With current computers, it is possible to treat two space dimensions and time for some problem types, three space dimensions for others, and three space dimensions plus time for a very limited set of problems. Speed increases of about a factor of 200 in this decade are needed to allow researchers to solve urgent multidimensional problems that are now intractable.

Resolution. Every region of space contains infinitely many points. Thus, the first step in modeling any natural phenomenon is to approximate the space with a finite set of zones, each of which requires a number of calculations. Increasing the number of zones means we can determine more completely and accurately what is happening in any environment, but the computational time grows very rapidly. For example, in a two-dimensional time-dependent model, the running time grows in proportion to the third power of the increase in resolution; increasing the number of zones by just a factor of 2 would increase the time to complete the problem by a factor of 8. Many complex problems now run up to 100 hours, so it is clear that resolution increases of even relatively small factors can overwhelm the capabilities of current supercomputers.

Physics. All computational models dealing with the frontiers of science and technology make simplifying assumptions about the laws of physics in order to keep the calculations from running too long. In some models, including just one additional physical effect can increase running time by a factor of 10. Faster supercomputers with much larger memories will permit researchers to solve problems that cannot now be economically solved.

Combination effects. Although dimensionality, resolution, and physics each have powerful effects on running time by themselves, the overall needs are derived from combinations of these effects. The highly complex problems now being studied in defense, space, energy, and science programs, as well as other programs in the national interest, require computational models with higher dimensionality, and with higher resolution, and with more physics.

Many of the applications for which these higher powered systems are needed have been described earlier in this section. A more complete description of additional requirements and new applications is provided in Appendix III.

Conclusion

Supercomputers have been regarded as powerful laboratory instruments which have significantly contributed (within the domains of the Federal Government) to the advancement of weapons, energy, and general scientific research. However, we have only begun to realize the potential impact that advanced computing could have upon our Nation's economy and security. Increasingly, supercomputers are also becoming recognized in industry and universities as a key to advancing engineering, manufacturing, commercial technology, and other areas of scientific research. Leadership in the application of much higher-performing supercomputer systems will ensure continued leadership and strength in this Nation's scientific and technological base. The Federal Government, industry, and universities will contribute to this leadership and strength and utilmately benefit from it.

THE COMPETITION

Supercomputers and Competition

Historically, supercomputers have been produced in the United States and they hold a unique place in American society. They are one of the fundamental tools needed for research in academic, government and private laboratories and have been critical factors in advancing the state-of-the-art in diverse fields such as aircraft design, oil exploration, and nuclear weapon development. As time passes, today's supercomputers, and parts of their architectural, software, algorithmic, peripheral, component, and packaging elements, become tomorrow's mainstream computer technology.

By the end of 1983 there will be approximately 75 supercomputers extant (60 in the U.S. and 15 outside the U.S.) having a total value of $700 million. Even an optimistic total of slightly more than $3 billion worldwide by the end of the decade is on the order of 1% of the total computer market. The reason this special product and this tiny market are so significant is that supercomputers are a key facet of quality computer systems. Even though the present market for supercomputers is an extremely limited one, recent product announcements have been made by Fujitsu (July 1982), Hitachi (August 1982) and NEC (until recently, known as Nippon Electric Company) (April 1983). When Hitachi announced on August 30, 1982, Haruyoshi Hatano, Director of Hitachi, stated: "This type of high speed machine is necessary so that we do not fall behind the Western countries in the technology development race," i.e., the Japanese thrust is intended to provide tools to advance their industrial and academic research at a rate intended to catch up or even forge ahead of their competition.

Competitive Threats

No nation's business philosophy is a static set of dogmas. The Japanese supercomputer thrust is a significant threat because it comes at a time when

- tactical & strategic business decisions in the U.S. are often dominated by a short term financial view
- relevant systematic long term efforts among Government, industry and universities are largely absent
- American productivity in related industries has been in a prolonged period of stagnation. This has produced a number of technological subcultures which have a negative self image.
- products resulting from Japanese long-lead-time efforts in high-speed high-density technology are available only from Japanese component vendors and are delivered first to Japanese system manufacturers.
- vertically integrated, broadly diversified Japanese organizations are competing with supercomputers produced by smaller specialized U.S. firms.

It is of no consequence whether the emergence of the Japanese supercomputers is a result of industrial policies in Japan or the absence of a sense of gestalt in American business philosophy. The United States must recognize the importance of supercomputer technology to developments at the leading edge in high technology and to the nation as a whole.

Impacts on the U.S. of Changing Leadership Role in Supercomputers

The U.S. is losing its worldwide monopoly in supercomputers and may have difficulty remaining competitive. Loss of leadership in supercomputer manufacture means the U.S. will lose immediate access to the fastest machines. This development would amount to a significant set-back in advancing state-of-the-art improvements in the design of weapons systems and aircraft. The severe negative consequences in national security involve several unattractive policy choices. These may mean using less capable or more costly technology to avoid reliance on foreign products; an erosion in the balance of payments situation; a loss of national pride, where persistent Japanese publicity has already had a negative effect far exceeding Japan's actual achievements. Even if the U.S. remains technically and economically competitive, the unique advantages we have enjoyed will be eroded, especially if foreign supercomputers find their way into the arsenals of our adversaries.

Supercomputer Projects

The following three tables list 1) existing supercomputers, 2) announced supercomputers and 3) supercomputer projects in the world. This tabulation employs only the few parameters usually contained in press-release-type information. This simple approach is not useful for technical comparison of these products and efforts. The tables are intended to illustrate which nations and companies are involved in supercomputer projects. The recently proposed DARPA project in the U.S. in not listed here since it is not considered to be a supercomputer project.* See Appendix I.

Table one shows supercomputer systems currently being installed. The first of the new generation of supercomputers, the CRAY 1S, is no longer being manufactured. It is not shown in the table although over fifty of them have been installed since 1976. It has been replaced by the Dual Processor CRAY XMP and a lower cost CRAY 1M. Approximately 15 CYBER 205s have been delivered. The first Hitachi S-810/20 is about to be delivered to the University of Tokyo. The first Fujitsu VP100 (a smaller version of the VP200) is to be delivered to the University of Nagoya in November 1983. It is to be used primarily for plasma physics research.

---

* "The emphasis in the DARPA program is in the area of symbolic processing rather than 'number crunchers'".

Statement by Dr. Robert S. Cooper, Director, Defense Advanced Research Projects Agency, before the Subcommittee on Investigations and Oversight, Subcommittee on Science, Research and Technology of the House Committee on Science and Technology, June 29, 1983; Pg. 6.

Table 2 shows new systems in design or early stages of prototype development. NEC is the latest to announce its entry, the SX 2, for delivery in 1985. The Cray 2 is the most developed and it is expected that several may be delivered before the end of 1984. However, it seems to be a prototype system for the newly designed Cray 3. The latter is intended to be the first Gallium Arsenide based supercomputer if this technology proves successful. The Cray 3, ETA GF 10, and Denelcor HEP 2 are promised for delivery in 1986. There is much risk in these programs, both timewise and performancewise. Each calls for the development of multi-processor systems containing at least eight CPUs. No multi-processing software is available for such systems and it is not clear that the 8-processor system will be 8 times more powerful than a single processor on a single problem. The time scale for the development of these systems has apparently been shortened in order to compete more favorably with Japanese products now in existence or planned. This speedup resulted primarily from the current intense supercomputer activity in the U.S. At the present time it is not known what plans Hitachi, Fujitsu, and NEC have for follow-on computers.

Table 3 shows the worldwide supercomputer projects that are Federally-supported. The intent is to develop components and architectures for the 1990's. Note that the U.S. is not among them.

TABLE 1

Current Supercomputers

Organization	Fujitsu	Hitachi	CRAY	CRAY	CDC
model	VP-200	S-810/20	1M	X-MP	205
announcement (or project start)	July 1982	Aug. 1982	Aug. 1982	Aug. 1982	June 1981
availability	Nov. 1983	Nov. 1983	Sept.1983	June 1983	Feb. 1982
architecture[3]	vector IBM compatible	vector IBM compatible	vector	vector multi-processor	vector
maximum performance (M FLOPS)1	500	630	250	630	400
maximum main memory size MB	256	256	32	32	64
Technology o Logic (gate delay)	ECL (350 ps)	ECL (350 ps)	ECL (SSI,.7ns)	ECL (.5-1ns)	ECL (700 ps)
o main memory	64K MOS static (55 ns)	16K MOS static (40 ns)	16K MOS (70ns)	4K ECL (25ns)	MOS (45ns)
Input/Output o No.paths (total)	32	32	19	19	16
o max data rate path (MB/sec)2	3	3	100	1,000	25
o max aggregate data rate (MB/sec)2	900	1,100	550	1,550	400

1 M FLOPS = millions of floating point operations per second
2 MB/Sec = millions of Bytes per second
3 all 64 bit words
4 derived from 6-12 times Cray 1
5 vector/options multiprocessor
NA : not available

TABLE 2

Supercomputers Now In Design*

Organization	CRAY	CRAY	ETA	Denelcor	NEC
model	2	3	GF10	HEP-2	SX-2
announcement (or project start)	none yet	none yet	Sept. 1983	May 1983	April 1983
availability	late 1984	1986	1986	1986	1985
architecture[3]	vector multi-processor		vector multi-processor	scalar[5]	vector
maximum performance (M FLOPS)[1]	1,000[4]	NA	10,000	4,000	1,300
maximum main memory size MB	NA	NA	256	4,000	256
Technology o Logic (gate delay)	ECL	ECL GaAs	CMOS (o.5-1ns)	ECL	NA
o main memory	MOS	MOS	MOS	ECL/MOS	64K MOS static
Input/Output o No.paths (total)	NA	NA	16	NA	32
o max data rate path (MB/sec)[2]			50-60		1.5
o max aggregate data rate (MB/sec)[2]	NA	NA	1,000	NA	1,350

1 M FLOPS = millions of floating point operations per second
2 MS/Sec = millions of Bytes per second
3 all 64 bit words
4 derived from 6-12 times Cray 1
5 vector/options multiprocessor
NA : not available

* Most are in stages of design and subject to change. The most developed is the Cray 2, but even in this case, not all parameters are known yet.

TABLE 3

National Supercomputer Projects

Organization	Japan (MITI Project)	France (Min. de la Defense DRET projects)		
model		ISIS	MARIANNE	MARISIS
announcement (or project start)	(Jan. 1981)	(1980) (BULL)	(1981) (SINTRA)	(1982)
availability	1989	1985/86	1986	1988
architecture[3]	vector multi-processor	vector	multi-processor	vector multi-processor
maximum performance (M FLOPS)1	10,000	50-100	low	200
maximum main memory size MB	1,000	256	NA	NA 800
Technology o Logic (gate delay) o main memory	JJ/GaAs (10-30 ps) 16K Ga As (10 ns)	NA	NA	NA
Input/Output o No.paths (total) o max data rate path (MB/sec)2	NA	NA	NA	NA
o max aggregate data rate (MB/sec)2	1,500	1,200	NA	NA

1 M FLOPS = millions of floating point operations per second
2 MS/Sec = millions of Bytes per second
3 all 64 bit words
4 derived from 6-12 times Cray 1
5 vector/options multiprocessor
NA : not available

SCIENTIFIC SUPERCOMPUTER TECHNOLOGIES:

ARCHITECTURE

Introduction

Supercomputers are used to solve computing problems beyond the capabilities of conventional machines. Their speed makes it possible to perform scientific and engineering simulations that are not otherwise practical because of time or cost limitations. Although the commercial market is now relatively small, a large user community is developing in such areas as structural analysis, computer-aided design, and advanced graphics--with future developments promising further market growth. Additionally, supercomputers traditionally have been the testbed for new components and new architectures. In the following paragraphs we briefly describe the major architectural classes and their relation to supercomputers.

The marketplace has always contained some computer users requiring significantly more capability than available with high volume mass produced machines. Machines built to satisfy those requirements, supercomputers, set the pace for the next generation of mass-produced machines. Today's supercomputers represent a major departure from past practice. Previous systems took advantage of improvements in component technology or made architectural changes that were transparent to the user. The reduced pace of component speed improvements has caused architects to look to radical new approaches involving, first, vector streaming and now multi-processing to achieve the speeds sought.

New architectures have been used in supercomputers, making possible speeds of up to 800 million floating point operations per second (MFLOPS). Traditional programs and algorithms achieve far lower speeds. More experience (which will come only with time and increased accessibility) is needed in order to extract better performance. Mechanisms must be created to make these new architectures accessible to a broad community of users. With the successes which will follow from experience, the cost-effectiveness of supercomputers, and the market, will grow.

The Von Neumann Paradigm

Virtually all computers constructed have used the Von Neumann architecture, employing a control unit, memory, and arithmetic unit. The control fetches from memory and examines program instructions sequentially and causes data to be fetched from memory, operated on in the arithmetic unit, and the result to be stored in memory. Each instruction operates on a single datum or combines a single datum with a value in the arithmetic unit.

Translation of computer languages into machine instructions is a straightforward matter because of the relatively simple architecture. Development of algorithms is similarly uncomplicated, with no need to account for the specialization found in other architectures.

Overlap and Other Speedup Tricks

To increase speed beyond that available at a given level of component technology while maintaining the appearance of a Von Neumann architecture required development of a number of special techniques. The common element is increased concurrency: several functions are performed at once so that the average execution rate is increased.

The first major example of concurrency is the overlap of arithmetic execution with the fetch of the next instruction. In this case, while the arithmetic unit is busy, the control unit can obtain the next instruction from memory. The effort to supply data at the rate possible with high-speed arithmetic units from slower memory prompted the development of interleaved banks of memories, which permit simultaneous accesses to data in different banks. Further, the availability of multiple arithmetic units permits simultaneous execution of arithmetic instructions.

Parallel and Vector Processors

Further concurrency was the next avenue to increase speed with a given component technology. Today's supercomputers eliminate the overhead of instruction processing associated with each datum in conventional machines. They permit a single instruction to specify that an operation be performed on multiple operands and differ from conventional systems primarily in the high bandwidth transfer of multiple data elements to and from the execution unit under the control of a single instruction.

When the execution unit operates simultaneously (in lockstep) on many data entities, as in the ILLIAC-IV, the machine is said to have an array architecture. The ILLIAC-IV contained an 8 x 8 array of floating point processor elements (PEs). Keeping most of the PEs busy most of the time is the major software and algorithmic challenge.

When the execution unit operates sequentially on sets of data, on an assembly line basis, the machine is termed a vector processor or pipeline processor. The CDC 205 and Cray 1 are examples of vector processors.

The speedup obtained is proportional to the number of stages of a vector processor or the number of PEs of a parallel processor. Scaling appears limited and is already a difficulty in vector processors, where arithmetic operations have already been decomposed into these basic steps.

The effectiveness, or throughput, of a supercomputer is a function of hardware capability, algorithmic complexity and software efficiency. In the case of array and vector processors, which demand a high degree of structure, it is necessary to develop algorithms to utilize the architecture properly. The potential speed of such machines is not realized when sections of programs cannot be put in a form (parallel or vector) amenable to the hardware. This usually results in significantly lower speed than peak performance.

A further problem is the lack of software for effective use of existing programs written for conventional machines. When such programs are run on array or parallel processors, they often do not take significant advantage of the processor's special capability.

Multiprocessors

Devoting hardware resources to the scheduling (structural) problems of machines offers the opportunity to keep the execution resources fully occupied. Multiprocessors, which execute several instruction streams simultaneously, each with its own data stream, may dynamically partition problems, transferring data to available execution units, and coordinating results as they become available. Specific examples include dataflow systems, where operations are performed as data become available and multiprocessor arrays with processors operating independently.

Models of the IBM 3081 will have two or four conventional systems linked through a common memory. Both Cray Research and Control Data are developing multiprocessors linking their large vector processors together in groups ranging from two to eight. Denelcor builds a four processor system.

We include the following examples for completeness even though they are not deemed supercomputers because of their lack of generality. Array processors (sometimes called convolvers), which attach to conventional machines, offer a low cost way of obtaining high computation rates for some problems such as signal processing and image reconstruction. Other special purpose implementations may find similar applications. The availability of VLSI components has led to speculation of "tailored" systems of special purpose subsystems built according to the needs of individual applications.

The algorithmic issues concerning whether there are or are not intrinsic problems in programming and efficiency or execution are not resolved today. It is generally agreed, however, that parallel processors can be scaled while vector processors cannot be scaled up much further.

Conclusion

New architecture for supercomputers will require wide exposure and experience before yielding full performance. It is not cost-effective to move codes to a new machine (nor, especially, a new architecture) unless a sizable speed increase is obtained. Thus it is important for the U.S. to capture the present market (by high performance gains) and lessen the attraction of marginally faster (2 times) machines announced by foreign competitors.

The key actions center upon widespread accessibility of machines produced in the U.S. This will create the success and experience base necessary for introduction of follow-on models. Among the mechanisms available to government are establishing access to its facilities, increasing NSF support marked for supercomputer access, supporting networks for the research community, and establishing supercomputer centers, for use by the scientific community.

Increased support for research aimed at the algorithmic and software issues of multiprocessing will decrease the time it takes to bring a new machine to full performance.

SCIENTIFIC SUPERCOMPUTER TECHNOLOGIES:

PERIPHERALS

In addition to the central processing unit, a computer system also contains other equipment, referred to as peripherals. These peripherals perform the following functions:

1. Entry of data and programs into the system.
2. Output of results from the system.
3. Storage of data within the system.

The ability of the supercomputer to process enormous amounts of data generates corresponding requirements for the input, output, and storage peripherals. Yet the thin supercomputer demand has strongly discouraged the development of peripheral equipment which can fulfill these requirements.

While supercomputers generate requirements for large amounts of data to be input, the normal practice is to record data on a storage media at the source and process it later. Consequently, this problem is more properly discussed in a following section on data storage.

Output from the computer (other than that directed to storage) can consist of printer listings, microfilm output, alphanumeric and graphical displays, and disposition to another computer for further processing or transmission to users over communications lines. The sheer volume of output can inundate a user with masses of data which must be converted into information. Graphical displays are one very useful method, but limited in application by the lack of three-dimensional representation.

The storage of data is the most important peripheral function, and is particularly critical to the supercomputer. The predominant ways to store data for a supercomputer today are:

1. On-line to the computer using magnetic disks.
2. Off-line using magnetic tape reels or cartridges.

Magnetic disks have the ability to store relatively modest amounts of data available to the computer in about 50 milliseconds, and which may be accessed randomly. Magnetic tape, on the other hand, can store much larger amounts off-line from the computer. However, the off-line nature of the storage means several minutes delay to retrieve the reels. Even then, the data may be accessed only in a serial mode. Mass storage systems utilizing cartridges and handling mechanisms are also available, but these suffer from low transfer rates.

A third type of device that has been sought for over 20 years is the archival storage device. This differs from magnetic disk and tape in that the data is not subject to change, must be stored for an extended period of time, and the volume exceeds the feasible capacity of disk and tape. Examples include tax records, results of experiments and seismic data. Several attempts based on photographic films, and other optical means, have been unsuccessful.

The difficulty of the supercomputer peripheral is clearly illustrated by citing an example of the disk situation. In 1972, a high performance disk was developed for attachment to the CDC 7600, the leading supercomputer of that time. Since then we have seen computers with 20-50 times the power of the 7600 appear. However, due to the lack of a market, the 1972-vintage disk is still used today on the two leading supercomputer models, and is, in fact, still the highest performance disk available.

Another factor which complicates the peripheral situation is the architectural change caused by the move to multiprocessor configurations. The traditional method of attaching peripherals is via a more or less direct path from the computer to the peripheral with one point of central control. The multiprocessor disturbs this straightforward arrangement and leads to the need for distributing control and data handling functions into separate processors. To handle the flow of data to multiple supercomputers requires input/output processors of considerable power with data handling capabilities greater than those of conventional large computers.

There is some significant technology under development which addresses these areas. In the magnetic disk area, a new technology called vertical recording is actively being pursued by several companies. The unfortunate aspect, from the supercomputer point of view, is that it does not appear that this technology will be initially applied to disks of direct interest to supercomputers. Instead, the first applications seem to be flexible disks for microcomputers, and to reduce the physical size and costs of disks for the micro, mini and medium scale computers. The logic driving this is the sheer size of the market at the small end: millions of units per year versus a few hundred for the supercomputer market.

Similarly, in the archival storage market, a number of companies are pursuing optical recording technology. The first products to appear, however, will be directed towards the mainline of the market.

There are specific peripheral equipment areas that require effort to insure balanced supercomputer systems.

- High performance disks using vertical recording.
- The development of suitable archival storage systems.
- The development of technology to meet the input/output requirements of the multiprocessor systems.

SCIENTIFIC SUPERCOMPUTER TECHNOLOGIES:

SEMICONDUCTORS

Electronic components are the heart of any computer system, with the performance of a given design governed primarily by component speed. Components have shown modest increases in performance and large increases in density over the past 15 years - a factor which has caused the supercomputer designer to turn to vector processing and large memories to achieve the current levels of performance.

The design time of a scientific supercomputer system ranges from five to ten years, depending upon the degree of new design required. The selection of component technology is made early in the logic design process, since the characteristics of the technology generates enormous effects on the design tools, packaging, and manufacturing processes.

The present generation of supercomputers, including the two major U:S. models and the three Japanese models, use the same silicon component technology referred to as bipolar. This particular technology has been used commonly for high performance applications for the past 10 years because of its speed. It will continue to be a candidate technology for such uses as improvements are made.

Technologies competitive to bipolar are being actively pursued by supercomputer designers and component suppliers. One of these is also based on silicon called CMOS while the other two utilize nonsilicon technologies. CMOS is the mainline silicon technology used in most components. Its appeal lies in the extremely high densities and low power consumption. ETA Systems * will be using this technology to implement their new supercomputer entrant.

Both the bipolar and CMOS silicon devices used in supercomputer applications are classified as uniquely designed components which will be manufactured in low volume at relatively high cost. These devices will not achieve the same circuit density as the standard high volume devices. This is a consequence of the desire to reduce chip design cost and partitioning problems which will reduce density. Various design techniques such as standard cell, gate array and full custom, will be used to design these components.

The customizable circuits such as gate array, PLA and other fixed cell approaches will find supercomputer use where the packaging technology provides a large number (200-500) of pins on the chip. The number of circuits per chip will remain relatively low (10,000-50,000),

Custom designed devices will take many forms. Semicustom design techniques, such as standard cell, will become prevalent for scientific supercomputer components due to the relatively high packing density achievable and the relatively low design cost. Full custom designed circuits will be employed sparingly, due to the very high design cost and the low production volume.

---

* ETA Systems was formed on Sept. 9, 1983, as a subsidiary of Control Data Corp. to develop supercomputers. The first supercomputers will be based on technologies pursued at CDC in designing the Cyber 2XX (now designated GF10).

Wafer-scale custom designs (Trilogy-type) may be ideal for scientific supercomputers and will be employed in specialized cases. However, it seems that this technology is more applicable to high volume commercial computers than to low volume scientific supercomputers. Therefore, one must ask if the wafer-scale, custom designed, commercial computer will be good scientific supercomputer, its low cost offsetting its low floating point performance.

Because device performance for supercomputers is of the essence, the device geometries and processes needed for these components will not generally be available on the so-called silicon foundry front-ends. Manufacturing costs will be relatively high because of the need for leading edge design rules and processes found only in high cost facilities. Through the end of the decade, design costs for custom designs will be in the neighborhood of $1 to $2 million per chip and manufacturing costs will be $500 to $2,000 per chip for chips of the complexity required by future supercomputers.

The VHSIC program of the Department of Defense is adding support to the technology of custom design silicon devices. This DoD program is providing R&D support for devices, circuits, systems and design automation. These technologies will find applicability in scientific supercomputer design over the next decade.

Two nonsilicon candidates which achieve some degree of use late in this decade are gallium arsenide and Josephson Junctions. These components must be viewed as possible components for scientific supercomputers of the 90's, not the 80's.

Gallium arsenide devices are the most likely to find use in scientific supercomputers during the 1990's. The major interest in gallium arsenide in the DoD stems from the requirements for high temperature and radiation tolerance. Today these devices require a large chip area and are at an SSI level of integration, at best. Whether the level of integration can be extended significantly is a subject of current research. Memory devices of 4K bit complexity and logic circuits of 1K complexity can be forecast. The packaging problems of these devices may restrict their use in scientific supercomputers to high density, exotic cooling system applications.

Cray Research Incorporated believes that in the 90's gallium arsenide will be the dominant material used to build computer logic for scientific supercomputers. With current gallium arsenide technology, a speedup of three is possible over silicon. That speedup can be tripled again with known techniques and tripled again by applying ballistic transport which works with gallium arsenide as well as with silicon.

It is worth noting that the Japanese under the direction of MITI have a significant gallium arsenide program under way in support of the Japanese supercomputer effort.

Josephson Junction scientific supercomputers have an uncertain future in the next decade. Although significant research has been conducted on this technology in the past, many of the leaders are de-emphasizing this effort. The problems to be overcome in the construction of even a small demonstration computer are significant. In addition to problems associated with the devices, there are system problems associated with shielding, cooling and I/O. There may be small systems designed, constructed and used late in this decade which serve as very special preprocessors for supercomputers.

The supercomputer places unique problems in packaging on the designer. High speed demands dense packaging of relatively high powered devices. Thus, the mechanical and thermal design of a supercomputer has been and continues to be a speed limiting factor. Continued advances in mechanical and thermal techniques will occur. Direct immersion cooling of devices will be employed along with special heat removal structures for each semiconductor device.

The reductions in device geometries already achieved are placing emphasis on interconnection systems. Many chips are now more than 50% metal on the top surface. Line widths and current densities must be reduced to achieve higher circuit densities and performance. Gallium arsenide offers the unique capability of optical communications on the chip which may prove useful. Multilevel interconnection systems must be improved and extended. Logic design methodologies should be developed which minimize interconnections rather than logic elements.

The role of design automation in the creation of components for scientific supercomputers should not be overlooked. The very long design cycle time mentioned in the introduction places unique demands on the design automation system to insure that the design time is not lengthened any further due to design errors. Thus, design automation can be considered a critical technology to the successful creation of scientific supercomputers.

The following areas deserve further attention:

- Gallium Arsenide Technology
- Josephson Junction Technology
- Packaging
- Electronic CAD

There are several new efforts aimed at increasing and coodinating semiconductor research in the U.S. The Semiconductor Research Cooperative (SRC) supports University research, using funds accessed from member corporations (a small fraction of semiconductor sales or use).

The Microelectronics & Computer Technology Corporation, (MCC)* also a consortium of computer and semiconductor manufacturers (accessed a fixed sum for each project supported) will carry out research in VLSI, CAD/CAM and Semiconductor packaging.

* See Appendix I.

SCIENTIFIC SUPERCOMPUTER TECHNOLOGIES:

SOFTWARE AND APPLICATIONS

Problem Statement

The next dramatic increase in speed in supercomputers will involve multiprocessor configurations and parallel structuring of the computations. The impact of explicit parallelism on software and applications will be enormous.

Parallel computing is a systems problem, not an architecture problem. Failure to recognize this salient fact will negate or substantially diminish the leverage of any program to develop economically usable very high performance computers based on parallel architectures. Exploitation of the current U.S. leadership position in software development capability is the strongest hope we have for reaping the economic benefits of applications of very high performance parallel computing.

Availability of algorithms and software which support efficient development and effective execution of programs implementing parallel computing constructs will set the pace for application of the next generation of supercomputers. Development of basic software for parallel architectures is intrinsically more difficult than development of basic software for sequentially structured systems. Thus software for parallel systems is likely to lag even further behind the hardware than is the case for conventional systems. This circumstance offers the U.S. a unique opportunity to "leapfrog" over Japanese and European competitors who have only partially recognized the importance of software and application research. The U.S. currently enjoys a substantial lead in software development capability. Exploitation of this lead can enable us to obtain the lead in economic exploitation of parallel very high performance computing almost regardless of leadership in parallel hardware and/or architectures. The role of Government funding for research and development of software and applications development is at least as critical as funding for research and development of parallel architectures. Most vendors state that costs for basic software exceed hardware costs. This reality must not be ignored. Failure to accelerate software and applications development will almost surely lead at best to a long delay in obtaining return on the investment in hardware and architecture.

The lowest level barrier to effective application of parallel computing is that there is a lack of detailed knowledge of parallel problem structuring and parallel models of computation upon which to ground development of applications and software. The software development problem is thus ultimately pushed back to applied theory, models of parallel computation and parallel algorithms and problem structures.

A two-faceted research and development thrust is needed.

1. vigorous exploitation of currently available parallel computing environments for software and applications development

2. Incorporation of application and software research and development with all parallel hardware and architecture research and development projects

New architectures can be evaluated only by their application. Failure to bring software along with hardware will delay evaluation of architectural concepts and thus delay progress in very high performance architectures. We should also have already learned that more usable machines will result from joint consideration of hardware, software and applications.

Problems for Research and Development

An effective base for the effective application of a parallel architectured next generation of supercomputers must be founded upon broadly applicable theories of parallel computation, a body of effective parallel algorithms for common computational tasks, effective software including usable programming systems for parallel programming and operating systems for parallel architecture and experience in building full application programs in parallel computing enviroments. These research topics can be described as follows:

Applicable models of parallel computation which can serve as design principles and tools for the development of algorithms, architectures and programs are fundamental.

There are simply not now available a body of applicable and well-characterized parallel algorithms for common computations tasks. Algorithms are the building blocks of programs. The need is to both establish techniques for parallel algorithm design and evaluation and to establish a repertoire of parallel algorithms applicable on a spectrum of architectures. The diversity of parallel architectures makes this task many times more difficult than the same task for sequential algorithms.

Programming systems which can effectively express processing involving large numbers of processs are simply not available. Attention must be paid to the human interface aspects of parallel programming. The complexity of mapping parallel computation structures of applications to parallel architectures brings computer science to unknown ground for the building of operating systems for parallel architectures.

The effectiveness of algorithms and software can be evaluated only by use. Substantial experiments generating parallel structured prototypes of major applications are a fundamental research requirement.

Government Role in Stimulating Research and Development

The research problems just defined are generic to the type of parallel computing intrinsic to the next generation of supercomputers.

Solutions for the problems posed will be a major stimulus to development and application of supercomputers in all discipline areas. There are at least two major actions which the Government can take which will accelerate parallel supercomputers through establishment of the basis for their applications.

Institutes should be established for parallel software and applications. They should be focused around a critical mass of researchers and both state-of-the-art commercially available and research prototype parallel architectures. The critical mass of researchers at each center would have one or more special areas of expertise such as programming systems, symbolic applications, and/or numerical applications. The facilities of the Institutes would be available either by a dedicated network or an available network to other research groups across the country. These Institutes will require budgets of from $1 million to $10 million per year to support both a critical mass of researchers and usable facilities.

Potential commercial applicability and the availability of research facilities in the Institutes just proposed will stimulate the research community to move into this problem area. Availability of funding to take advantage of this opportunity to encourage national goals is vital. It is particularly important that small business have the opportunity to compete for these funds. Entrepreneurial small business has become a major force in innovation of software for small computers. Its role in application of large scale computing may be very significant.

APPENDIX I

Definition of Fifth Generation Computers and Distinction From Supercomputers

The term "fifth" generation has been assigned by the Japanese, somewhat arbitrarily, to indicate computer systems of the 1990's which possess attributes desirable from the viewpoint of the Japanese society. This use of the numeral "fifth" implies a definition of computer "generation" in terms of time - a generation being approximately ten years. According to this categorization, the first generation was in the 1950's and, hence, the fifth generation in the 1990's. Actually, the terms first and second generation were quite well-defined by progress which began in the late 1940's and resulted in transistorized computers in the late 1950's and early 1960's. Somewhere in the mid-1960's the term "generation" came to have dual meaning. From one point-of-view, it denoted time (as in the Japanese use of the term). From another point-of-view, it denoted progress (having to do with significance). Recent Japanese use of the term in the time sense seems to have clarified the use of the term for future discussions. However, the term has not been consistently used in this sense during the 1960's and 1970's.

Outside of Japan the term Fifth Generation has been used loosely. In the United States it is most often identified with artificial intelligence capabilities (one key aspect of the Japanese Fifth Generation goals relates to improvements in this area). Press articles (e.g. Newsweek cover story July 4, 1983) have usurped the term "supercomputer" to mean a computer with extraordinary reasoning and intelligence properties. This use mixes a number of advanced computer properties in to one melting pot called "super"-computer.

There are a number of reasons for the semantic difficulties with description of advanced computer systems intended for application a decade or so in the future. The first and most fundamental reason is that the scientific community itself has no complete or simple description of a Fifth Generation Computer System and no such system exists today. The same ambiguity is not the case for supercomputers. Supercomputers exist today, their characteristics are completely describable, and reasonably simple summaries of these characteristics may be made. The second reason for the potential confusion lies in the response of western nations to the Japanese Fifth Generation Computer Systems Conference. This has led to a number of diverse activities to be described as Fifth Generation Computer activities in the sense that they are intended for application in the 1990 time frame. Hence, a comparison of Fifth Generation Computer Systems to supercomputer systems is effectively a comparison of apples to oranges.

APPENDIX II

BIOGRAPHIES

DR. SIDNEY FERNBACH, Chairman, graduated from Temple University in 1940 with an A.M. degree in Theoretical Physics; Computer Science. In 1952, he received his Ph.D. in Physics from the University of California. His professional experience is as follows: Physicist, Frankfort Arsenal, 1940-43; Assistant Instructor, University of Pennsylvania, 1943-44; Assistant, University of California, 1946-48; Physicist, Lawrence Radiation Lab, 1948-51; Member staff, Stanford University, 1951-52; head Theoretical Division, 1958-68, Head Computer Department, 1962-1977, Dept. Associate Director, Scientific Support, Lawrence Livermore Laboratory, 1977. Concurrent Positions: Member of the Board of Directors, Advan. Memory Systs., Sunnyvale, California, 1969; Consultant, U.S. Dept. of Defense, 1971; Consultant, Computer Systems Technology Adv. Comt, Domestic & Iternational Business Administration Bur East West Trade, U.S. Dept. Com. 1972; Member Steering Committee, Computer Science & Engineering Research Study, NSF, 1974; Member, Fellow American Physics Society; Asn. Computer Mach., Soc. Indust. & Appl. Math, Math Asn. Am. Inst. Electrical & Electronics Engineers, Research, Neutron physics, cosmic ray shower theory, computation.

DR. ALFRED E. BRENNER received a B.S. in Physics from the Massachusetts Institute of Technology in 1953 and a Ph.D. (physics)in 1958. His professional experience includes: Ford Foundation Fellowship - Physics, European Organization for Nuclear Research, 1958-59; Instrumentation and research fellowship, Harvard University, 1959-62, Assistant Professor, 1962-66, Sr. Research Associate, 1966-70; Physicist and Director - Computers, Fermi National Accelerator Laboratory, 1970; Concurrent Positions: Consultant, Department of Energy, 1972 and Summagraphics Corporation, 1973; Member, Com Dept. Computer Systems Technical Advisory Committee, 1974; Member: AAAS; American Physics Society; Association Computer Machinery. Research: Experimental high energy physics; computers and data processing.

DR. JAMES C. BROWNE received a B.A. Degree from Hendrix College, Conway, Arkansas, in 1956 and the Ph.D. degree in physical chemistry from the University of Texas at Austin in 1960.

He worked from 1960 to 1964 as an Assistant Professor of Physics, University of Texas at Austin, carrying out quantum mechanical calculations of small molecules, and doing research and development work on operating systems. He was, from 1964 to 1965, at the Queen's University of Belfast, Belfast, Northern Ireland, on a National Science Foundation Postdoctoral Fellowship. He was, from 1965 to 1968, Professor of Computer Science and Director of the Computation Laboratory at the Queen's University of Belfast. In 1968 he rejoined the University of Texas at Austin as Professor of Computer Sciences and Physics, and served as Chairman of the Department of Computer Sciences in 1968/69 and from September 1971 until January 1975. His research interests include operating systems, and theoretical computations of molecular properties. He has published about 125 refereed papers on technical subjects.

Dr. Browne is a Fellow of the American Physical Society and the British Computer Society, and a member of the Association or Computing Machinery.

MR. HARVEY CRAGON is a Texas Instruments Senior Fellow and has been responsible for the design and development of several digital computers during the past 25 years, including the Texas Instruments Advanced Scientific Computer used at GFDL Princeton University, the Naval Research Laboratory, and Texas Instruments. His current interests are in the field of computer architectures intended for efficient execution of artificial intelligence programs. He is also developing expert systems to assist in the architectural design of computers and microprocessors. The goal of this work is to have the ability to create error free specifications. He has been an adjunct professor of Computer Science at Southern Methodist University since 1977. He has a BSEE from Louisiana Polytechnic Institute and completed graduate studies at UCLA and SMU. He holds eight U.S. patents and is a member of the National Academy of Engineering, Association for Computing Machinery, Fellow of the IEEE, a member of the IEEE Computer Society and a member of the Charles Babbage Institute. He has been program chairman for many international conferences and has served on numerous panels and committees.

DR. JAMES F. DECKER received his B.S. from Union College in 1962, his M.S. in Physics from Yale University in 1963, and his Ph.D. in Physics from Yale in 1967. He was a member of the Technical Staff at Bell Telephone Laboratories from 1967-1975. (Research in plasma physics, integrated circuit development) He joined the Atomic Energy Commission in 1973 to work in the magnetic fusion program and has held the following positions in AEC, ERDA and DOE: Branch Chief for Experimental Plasma Research, Division Director for Applied Plasma Physics, and Special Assistant to the Director of the Office of Energy Research (present position).

DR. TSE-YUN FENG (S'61-M'67-SM'75-F'80) received the B.S. degree from the National Taiwan University, Taipei, Taiwan, the M.S. degree from Oklahoma State University, Stillwater, and the Ph.D. degree from the University of Michigan, Ann Arbor, all in electrical engineering.

He was on the Faculty of the Dept. of Electrical and Computer Engineering, Syracuse University, Syracuse, NY, from 1967-1975, and

Wayne State University, Detroit, MI, from 1975 to 1979. He is now a Professor in the Dept. of Computer and Information Science, Ohio State University, Columbus, OH. He has extensive technical publications in the areas of associative processing, parallel and concurrent processors, computer architecture, switching theory, and logic design, and has edited four Sagamore Computer Conference Proceedings (including *Parallel Processing*, Springer-Verlag).

DR. FRANKLIN F. KUO received his B.S. in electrical engineering from the University of Illinois, Urbana, ILL, in 1955 and his M.S. in 1956. Presently he is the Senior Scientific Advisor Executive Director at the Center for Intelligent Computer Systems Computer Science and Technology Division at SRI International. His specialized professional competence is in the areas of: Computer-communication networks, data communications, satellite data networks and military information systems. He has had 70 papers and 8 books published.

Professional Experience: Prof. of Electrical Engineering, University of Hawaii 1966-82; Director of Information Systems, Office of the Assistant Secretary of Defense. (C [3] I) 1976-77; Liaison Scientist, U.S. Office of Naval Research, London, England, 1971-72; Member of Technical Staff, Bell Laboratories 1960-66; Assistant Prof. of Electrical Engineering, Polytechnic Institute of Brooklyn, 1958-60;

Other Professional Experience: Member of the CNO Executive Panel, Subpanel on Force Enhancement, 1980 - present; Member of the COSINE (Computer Science in Electrical Engineering) Committee, National Academy of Engineering, 1966-72; Consultant, Lawrence Livermore Laboratory, 1966-71; Consulting Editor, Prentice-Hall Series on Computer Applications in Electrical Engineering 1967-present.

DR. LAWRENCE A. LEE currently is a Computer Scientist in the Atmospheric Sciences Division of the National Science Foundation. He is the Program Coordinator for Computing to the National Center for Atmospheric Research in Boulder, CO. He came to the Foundation in 1976 as a Mathematical Statistician and in 1978 was given the assignment of Computer Specialist. Previous to his entrance in the Federal Government (1969-1976) he was a Professor of Mathematics at the U.S. Naval Academy in Annapolis, MD. He was a Graduate Teaching Assistant in Mathematics at George Washington University, 1965-69.

Dr. Lee received his Ph.D. in Mathematics in 1975 from the George Washington University.

Some of his major committee work has been with the NSF Management Information Advisory Group, and the NSF Task Group on Scientific and Engineering Needs in Computing Resources. He was an organizer for the DOD/NSF Panel on Large Scale Computing for Science and Engineering in 1982. He served on the NASA-Ames User Steering Committee for the Numerical Aerodynamic Simulation from 1980-83. He was an NSF Observer from 1980-83 on the National Center for Atmospheric Research Scientific Computing Division Advisory Committee and was Chairperson of the NSF/AAEO Directorate Task Group on Office Automation in 1981.

MAJOR CARL EDWARD OLIVER received his B.S. in Mathematics from the University of Alabama at Tuscaloosa, in 1965; M.A. in 1967, and Ph.D. in 1969 (math). His professional experience is as follows: Assistant Professor math, U.S. Air Force Institute of Technology, 1969-74; Mathematician, Weapons Lab, U.S. Air Force, 1974-78; Program Manager, Numerical Analysis, U.S. Air Force Office of Scientific Research, 1976. Concurrent Positions: Lecturer, Ohio State University, 1970-71; Wright State University, 1971-74; and New Mexico Highlands University, 1974-78. He is a member of the American Math Society and the Math Association of America. He has done research in the following areas: interrelations and applications of the concepts within integration theory; numerical linear algebra, especially sparse matrix theory.

DR. JOHN P. RIGANATI has been Division Chief at the Institute for Computer Sciences and Technology (ICST) of the National Bureau of Standards (NBS) since 1979. His Computer Systems Components Division addresses computer data communication, advanced computer architectures, data interchange media and mechanisms, and standards for large and small scale computer interfaces. Laboratory activities include magnetic and optical storage media, speech I/O, data encryption and integrity, high-speed and local networking, data communication protocol verification, and a highly survivable network for military applications. Since Oct. 1981, when he was a panelist at the Fifth Generation Computer Systems Conference in Tokyo, he has focused on technical and strategic developments in Japan and has been active in formulating appropriate U.S. responses.

Professional Experience: Chief Scientist and Manager at Rockwell International's Communication Switching Systems Div. in Anaheim, CA, 1976-79; Rockwell International's Electronics Research Div., 1969-76.

Dr. Riganati has over 37 papers, reports and book chapters and holds seven patents.

DR. PAUL SCHNECK holds the Ph.D. in Computer Science from the Courant Institute at New York University, and the B.S. and M.S. from Columbia University's School of Engineering. He is a senior member of the IEEE and a member of the Association for Computing Machinery and the British Computer Society.

Currently he is Leader of the Information Sciences Division of the Office of Naval Research. Previously he was Assistant Director for Computer and Information Science at the Goddard Space Flight Center (GSFC) NASA, 1981-83; headed an Agency-wide team to plan and implement a NASA-wide administrative computing system; was a member of the Program Review Board for the Numerical Aerodynamic Simulation Project; remains Chairman of the NASA Deputy Administrator's Computer Advisory Group; served as Chief of the Information Extraction Division at GSFC, 1980-81; Director of the NASA Summer Study in Computer Science, 1981.

DR. DANIEL L. SLOTNICK is currently employed as Prof. of Computer Science at the University of Illinois, Urbana, ILL. He is an independent Consultant to domestic and foreign computer-related corporations and United States Government agencies including: NASA, U.S. Dept. of Defense, Standard Oil of Indiana, IBM, Atlantic Refining Co., Intel and Burroughs (since 1957).

Professional Experience: Director (Principal Investigator) of ILLIAC IV Computer Project, 1965-74; Director, Center for Advanced Computation, 1970-73; Manager, Computer and Data Systems Group of the Aerospace Division of Westinghouse Electric Corp., 1960-65; Advisory Engineer, Research and Advanced Systems Div., International Business Machines Corp., 1957-60; Research Assistant, Nonlinear Mechanics Group, Princeton University, 1956-57; Analyst, Electronic Computer Project, Institute for Advanced Study, 1952-54;

Other Professional Experience: Chairman, NASA Computer Directions Planning Group, 1981; Lecturer to National Security Agency professional staff, 1980; Representative IFIP (International Federation of Information Processing Societies) Working Group 10.3-Computer Design, 1978-present; Member Technical Advisory Committee Northeast Illinois Planning Commission, 1973-75.

Dr. Slotnick has over 200 internal design documents and reports at IBM, Westinghouse and at the University of Illinois while in charge of the ILLIAC IV Project.

MR. HUGH F. WALSH is the Manager of Systems Planning & Analysis in the Scientific & Engineering Processor Products area of the IBM Data Systems Division, Kingston, N.Y. He has been with IBM since 1955 in Technical Marketing, Education and Systems Planning. Since 1960 he has been involved in the planning of IBM's large systems, including the System 360 Models 91 and 95 and the System 370 Model 195.

Prior to joining IBM, he was Associate Professor of Physics at Siena College, Loudonville, N.Y.

Mr. Walsh holds a Bachelor of Science degree in Physics from Siena College, Loudonville, N.Y. and a Masters of Science in Physics from Rensselaer Polytechnic Institute, Troy, N.Y.

ADVISORY MEMBERS

MR. THEODORE H. BONN was named Director of the Computer Research Laboratory, the Sperry Research Center, Sudbury, MA, in 1970, and is now with the Sperry Defense Electronic Systems Operation. He is a member of the Board of Directors of the IEEE, representing the IEEE Computer Society. During 1980 and 1981 he was Vice President, Publications of the IEEE and a member of the Executive Committee.

Mr. Bonn started in the computer field with the Eckert Mauchly Computer Corp. in 1947. He participated in the development of the Univac I, LARC (the first supercomputer), Univac Solid State 80/90 and numerous peripherals and components. Mr. Bonn was Chief Engineer for Research and Peripherals of Univac. He is a Fellow of the IEEE.

He received a B.S. degree (1943) and a M.S. degree (1947) both in Electrical Engineering from the University of Pennsylvania.

DR. JACK LUBOWSKY holds a BEE from The City College of New York and an MSEE and Ph.D. (EE) from the Polytechnic Institute of Brooklyn. He is Director of the Scientific Computing Center at the State University of New York's Downstate Medical Center, an Associate Professor of Neurology in the College of Medicine and an Associate Professor in the Graduate School.

Dr. Lubowsky has previously worked at Airborne Instuments Laboratory on biomedical instrumentation and at Brookhaven National Laboratories on high speed nuclear instrumentation. During 1983 he was selected to serve as a AAAS/IEEE Congressional Science Fellow for the Senate Subcommittee on Energy, Nuclear Proliferation and Government Processes under Senator John Glenn.

He is a senior member of IEEE and has served in a variety of capacities including chairman of the Princeton Section.

Dr. Lubowsky has worked and consulted extensively in the application of computers and instrumentation to research, particularly biomedical research and has over 18 publications.

MR. DOYCE E. SATTERFIELD's professional career began in 1964 with the NASA's Marshall Space Flight Center where he was responsible for the design and development of both manual and automatic ground support equipment for the control and test of the S-1C stage of the Saturn V moon rocket program. Mr. Satterfield was an active member in the control center for the test and evaluation of the first 11 static test firings. Following this, he accepted a position with the Advanced Sensors Laboratory of the Army Missile Command. While there, Mr. Satterfield was responsible for the design and development of sensor and data processing systems for anti-tank missile programs. In 1971, he joined the Department of Army's Ballistic Missile Defense Advanced Technology Center.

With the Advanced Technology Center, Mr. Satterfield has been assigned to three major research areas. As a member of the Technology Analysis Directorate, he was responsible for the synthesis and analysis of advanced ballistic missile defense systems and the evaluation of related technology. As a project manager for the Radar Directorate, he directed the research and development of advanced radar components concentrating on large phased array radar systems. As a senior manager in the Data Processing Directorate, he directed the research and development of advanced hardware and advanced software for both centralized and distributed data processing systems. Currently, Mr. Satterfield is a senior management employee of the Ballistic Missile Defense Program Office, developing short-, mid-, and long-range program goals, plans, and objectives and providing advice/recommendations to the BMD Program Manager.

APPENDIX III

Some Existing Applications For Which More Powerful Supercomputers Are Required

Weapons Program. Since the beginnings of the Manhattan project, computation has played a central role in the understanding and designing of nuclear weapons. As the required designs have become more demanding and more difficult, the computational models necessary to produce these designs have become increasingly complex and have required more computing power. New developments in the areas of hydrodynamics and transport theory promise to provide the tools necessary to deal with this complexity. These new developments have one property in common; they require considerably more computational power than the techniques they replace. Unless weapons designers can acquire, within a few years, significantly more powerful supercomputers, their ability to develop and use the new computational models will be severely constrained.

Magnetic Fusion Energy. The central focus of magnetic fusion theory is the behavior of plasma as it interacts with magnetic fields. The growing body of physics theory that pertains to this interaction is extremely complex. Supercomputers have permitted scientists to simulate plasma actions. However, these models need to be more accurate and, in general, more applicable. This will require a computer with about 200 times the CPU speed of the Cray-1 and several hundred million words of memory in order to keep the running times down to a few tens of hours.

Inertial Confinement Fusion. In this research area, pellets of fusion fuel are compressed with intense pulses of energy. The goal is to elucidate some of the physics associated with nuclear explosives and to study its prospects for commercial fusion energy. For the past several years it has become apparent that the available computing resources limit the calculational capabilities of this program. Considerable progress has been made possible through the use of the supercomputer, which runs ICF codes some 10 times faster than they ran on the earlier computers, but even with this faster computer, current calculations consume up to 100 hours of computer time each. It is vital that the ICF program be able to obtain at least two orders of magnitude increase in computing power in the near future to satisfy these needs.

Reactor Safety. Computer modeling of reactor accidents is an essential aspect of the commercialization of nuclear power in the United States. Because full-scale experiments simulating severe accidents are not practical, the accident behavior of these plants including the adequacy of emergency systems must be evaluated by applying complex computer codes. These codes are very large (about 50,000 source statements) and run times are so long that they often require a week or more to complete a single case. Much faster computers are needed so these advanced codes can be used to resolve urgent safety issues in a timely manner.

Space Science Program. The influence and utilization of supercomputers on space science has greatly enhanced the ability to solve difficult and complicated questions about the nature of our universe. Since the data acquired from instruments on spacecraft are not only complicated in form but also voluminous, this science is wholly dependent on computers, especially supercomputers. Each data stream acquired from a space science mission must be formatted, reduced, catalogued, archived, and finally disseminated to scientific investigators for analyzing and comparing to model output. The data volume currently being acquired are in excess of 10(14) bits per year. This may increase by two orders of magnitude by 1990. Without more powerful supercomputer systems, processing and scientific utilization of space-derived data will substantially degrade.

Weather Service. The National Weather Service provides weather forecasts to the general public, issues warnings against tornadoes, hurricanes, floods, and other atmospheric and hydrologic hazards, and provides a broad array of special services ot aeronautical, maritime, astronautic, agricultural, and other weather sensitive activities. It depends on supercomputers for the collecting and archiving of meteorological data and for the developing of short, medium, and long-range forecasts. Much greater performing supercomputers are needed for improving forecasts and for processing and interactive analysis of the vast amounts of real-time meteorological data.

Geophysical Sciences. The atmospheric and oceanic sciences involve research that require extensive large-scale computational support. The nature of atmospheric and oceanic phenomena and the physical and temporal scales involved results in extensive data archives containing observations of those phenomena. New and more accurate models are needed. Examples are: global atmospheric-oceanic climate, ocean circulation, thunderstorm and tornado, solar cycle, turbulence, and atmospheric chemical dynamics. In earth sciences supercomputers are used for developing numerical models of subsurface systems, and the processing and analysis of geophysical data. Current research in these areas addresses ever more complex systems, emphasizing a greater degree of realism, completeness, and detail. These kinds of calculations need the fastest computer available at any given time. Thus, it is estimated that the geophysical sciences need several orders of magnitude increase in supercomputer capability (speed and memory) in the next several years.

Physical and Mathematical Sciences. In many fields of theoretical physics faster supercomputers with much larger memories will provide more realistic opportunities for numerically attacking difficult problems in quark confinement, turbulent behavior in classical fluids, quantum fields and plasmas, atomic and nuclear scattering and structure, and the behavior of fully relativistic systems such as neutron stars and black holes. Theoretical chemistry tends to be extremely computer intensive. Nevertheless, there are many important chemical problems whose solution effectively exceeds the capability of existing supercomputers. Many of these problems are biologically oriented because of the very large systems that arise there. One example would be the computation of the energy levels of a protein-sized molecule. There are also important problems in combustion kinetics which need more accurate calculations and thus require an increase of two or three orders of magnitude in computing power.

The role of the supercomputer is significant, and in fact crucial, in several important areas of mathematical sciences. For example, there are many areas of applied mathematics in which one needs to find solutions to partial differential equations or numerical solutions of large systems of equations. Significantly greater supercomputer performance is essential if mathematical scientists are to be able to continue to penetrate at the cutting edge. Fluid dynamics, mathematical modeling of large systems, cryptography, Monte-Carlo methods -- all require this. Of equal importance is the central role mathematics plays in the design of supercomputers, i.e., the development and application of efficient algorithms for parallel processors, numerical analysis methods and methodologies, etc...

Engineering Support. In many diverse industrial and Government activities, some engineering activities have come to rely quite heavily on supercomputers in solving problems supporting design efforts for future products or for increased productivity. Currently, taking advantage of such support are the automotive, aircraft and aerospace, petroleum, and semi-conductor industries. Similarly, at every major research laboratory, both Government and industrial, numerous engineering support aids require the use of supercomputers.

For example, the finite element method, which had its origins in structural mechanics, is being used increasingly in many diverse technical areas stretching into such areas as accelerator magnet design and hydraulics, hydrology and water resources. This technique has increasingly stringent requirements for high computing speed and large storage. There are may areas of structural mechanics which presently utilize and could greatly benefit from the increased availability of present and future generations of supercomputers.

Other areas, currently supported by supercomputers, include the design of the most advanced large-scale integrated (VLSI) circuits in the semi-conductor industry. Here computer simulation and auto-routing of the interconnections between circuit elements is crucial to the success fo such complex engineering endeavors. And, as the complexity and density of these circuits increase, the need for, indeed the possibility to do this job at all, will depend upon the avialability of faster supercomputers.

The current exploding growth in this use of computer driven graphics to support such diverse activities as the study of hydrodynamics and the production of highly popular animated motion pictures is paced primarily by the availability of supercomputing capability. Here the use of such tools can affect almost every industry, both new and old. The pace is dictated by the availability of the supercomputer cycles and the tools to use them.

New Applications

There are many areas of science and industry only now beginning to take advantage of the special capabilities that supercomputers make available. There are many more nascent applications, some not yet adequately perceived, that will surely take advantage of supercomputers to bring new levels of productivity to industry, and new depths of scientific discovery and engineering developments not heretofore feasible. Those industries that do not learn to do this will certainly be surpassed and made obsolete by those that do. Within any given industry, those firms that do learn how to properly use these tools will surely gain an economic advantage over their competitors. Thus, for example, computer aided design (CAD) tools are already in extensive use by industries as diverse as the architectural/engineering, automotive, aircraft and aerospace, electronics, piping, entertainment and garment industries. Competitors who do not learn to use these tools will certainly lose out. Although currently only a small fraction of CAD activity requires the use of computers as powerful as supercomputers, that aspect of this requirement will grow as industry learns its advantages. The breadth of applications will also grow as new supercomputer architectures are developed and produced as products.

Some applications just beginning to use supercomputers are:

High Energy Physics. The experimental high energy physics programs have used extensively both mini-computers for data acquisition and large computers for the analysis of that data in their search for fundamental understanding of the accelerators and larger and more complex detectors with which to perform these important experiments, it is mandatory to have available the largest supercomputers with the most advanced architectures as possible. These are required to test and optimize the design of both the accelerators and detectors, with simulation techniques, before actually committing to the costly enterprise of constructing the devices. Progress in this field depends crucially on having that supercomputing power available.

Biology. While small computers for controlling experimental setups and performing routine calculations are common, more and more emphasis is being placed on computers for more sophisticated, extensive and complex investigations. The three dimensional visualization of macromolecules has become essential in the study of macromolecular interactions. If a supercomputer were to be operated in conjunction with a molecular graphics system, one could carry out molecular modeling in real time. On the cellular scale, the three dimensional reconstruciton of a cellular structure form serial sections is necessary to understand that structure and its functioning. This is a large and complex problem that requires more capable supercomputers. Thus, the need for more powerful supercomputers in the biological area is real and in some cases critical. It is unlikely that the more difficult and important problems in biological macromolecular chemistry will be solved without them.

Social and Economic Science. Experimental research is one of the most active research areas in the social sciences. Large computer systems have been used in some of these experiments to control experiments and monitor their results in ways that would no otherwise be possible. Access to more powerful supercomputers would permit more ambitious experimental testing of economic, political, and social theories. A new generation of economic models is emerging. These models could dramatically improve the rigor and accuracy of economic analysis. But some fo the features of these models such as the large number of interdependent variables, the complicated optimizing solution methods, and the large scale of some of th emodesl plae a significant burden on today's supercoputer facilities.

Behavioral Sciences. Cognitive scientists are intensive computer users. Supercomputers are the next step for use in research in this field. If faster supercomputers became available, it is likely that within a 3-4 year time span, modeling of humans within a particular subject area, such as high school algebra, would be feasible. This activity could have the potential of revolutionizing cognitive science because it will provide the opportunity to deal with real world problems - particularly educational methodologies.

APPENDIX IV

SPEAKERS OR PARTICIPANTS IN SPECIAL MEETINGS

Ms. Prudence Adler, Office of Technology Assessment/U.S. Congress

Mr. Justin Bloom, Former Counselor for Scientific &
Technological Affairs, U.S. Embassy
Tokyo & London

Mr. Robert A. Burmeister, Hewlett-Packard

Admiral Bobby R. Inman, Microelectronics & Computer Technology
Corporation (MCC), U.S.N.-retired

Dr. Robert E. Kahn, Defense Advanced Research Projects Agency (DARPA)
Information Processing Techniques Office (IPTO)

Dr. John Killeen, Lawrence Livermore National Laboratory

Mr. Robert A. Koenig, ETA Systems, Inc.

Dr. David J. Kuck, University of Illinois

Dr. David Nowak, Lawrence Livermore National Laboratory

Mr. James L. Previte, Griffith Air Force Base/Rome, N.Y.

Mr. Kermith Speiermann, Chief Scientist, NSA

Mr. Martin T. Storma, Cray Research, Inc.

Mr. Lloyd M. Thorndyke, Control Data Corp.- ETA Systems

Mr. Ronald Ward, SOHIO Petroleum Company

Dr. Hans H. Zappe, IBM Research

REPORT TO
THE FEDERAL COORDINATING COUNCIL ON SCIENCE, ENGINEERING AND TECHNOLOGY SUPERCOMPUTER PANEL
ON
RECOMMENDED GOVERNMENT ACTIONS TO PROVIDE ACCESS TO SUPERCOMPUTERS

I. Executive Summary

A. Background

In January 1983, a Federal Coordinating Council on Science, Engineering and Technology (FCCSET) Panel on Supercomputers was formed to examine what, if anything, the U.S. Government should do to stimulate the use and development of supercomputers. As a result of the FCCSET panel discussions, three interagency working groups were formed. These groups are to examine the following supercomputer issues:

1. What should the Government do to ensure that the U.S. retains its lead in supercomputers? This group is named the Procurement Group.
2. What should the Government do to make supercomputers available to more researchers, particularly in universities? This group is named the Access Group.
3. A third group was formed to provide coordination among Government agencies funding research contributing to the technology base.

This report describes the findings and recommendations of the Access Group. The Access Group consists of the following members:

James F. Decker, DOE (Chairman)
Lawrence Lee, NSF
Robert Kahn, DOD
Kenneth Walgren, NASA
Richard Dunlavey, DOC
George Rogers, DCI

The Access Group was concerned only with making policy recommendations on access for scientists and engineers to supercomputers. Supercomputers are the most powerful computers available at any given time for doing scientific calculations. Future Fifth Generation machines for artificial intelligence applications were not considered. The Group relied upon previous studies for documentation that a significant need for access exists and it did not attempt to quantitatively assess the access requirements in the short time available. Further, the Group did not examine research and development needs for supercomputing, nor did it consider other types of computing.

B. Findings

1. The requirements for large-scale computing to solve many of the nations important research and development problems have been well documented.

2. At the present time, most federally-funded supercomputers are used by individual laboratories or by single programs. Very few researchers, particularly in universities, have access to these machines and, as a result, are unable to address many forefront research problems.

3. The need for more supercomputing in the U.S. research community cannot be met simply by wider access to existing federally funded computer resources in government laboratories. Many federally funded facilities are operating at or near full capacity, or plan to use available capacity for future programmatic growth. Facilities that now support research in weapons development or nuclear energy, for example, would have to be totally restructured for unclassified users to share computing resources. The three university supercomputer centers are underutilized. These universities have had difficulty in obtaining sufficient resources to provide the necessary, broad, user friendly services. Commercial service bureaus do have available computer time to sell.

4. The programmatic need to provide access to researchers geographically dispersed across the country, and economics dictate serving users remotely by communications networks from a limited number of supercomputer centers.

5. High quality service which maximizes the productivity of researchers should be a prime consideration when establishing supercomputer services for remote users. Several years will be required to obtain high quality service from a newly established center.

6. There are two federally supported supercomputer centers that have a relatively long history of providing good service to large numbers of remote users: the Magnetic Fusion Energy (MFE) Computer Center at the Lawrence Livermore National Laboratory, and the National Center for Atmospheric Research (NCAR) Computer Center in Boulder, Colorado.

7. There are insufficient opportunities for training students in the physical sciences and engineering in the use of supercomputers and for training computer science students in the operation of supercomputer systems.

C. Recommendations

The Government should take action to provide access to supercomputers for government funded researchers in government laboratories, universities, and other institutions who have large computational problems requiring the use of supercomputers. Supercomputers are critical in solving the complex problems vital to our security, economic, and scientific interest. There is clear evidence that the lack of access to supercomputers limits progress in many aspects of science and engineering.

The Working Groups's recommendations were developed to help solve this national problem. Recommendations are made on near-term approaches to providing access to large numbers of users quickly and also for the longer term. The Group tried to identify those features of the MFE and NCAR operations (described in Appendix A) which have made them successful in providing high quality services to remote users. Recognition of those features was a key factor in developing many of the recommendations.

A continuing interagency group should be designated to provide coordination among agencies to help ensure appropriate sharing of resources, to avoid duplication of effort and to maintain a dialog on policy and technical issues.

To meet immediate needs, each federal agency supporting scientific and engineering research should:

1. if not already accomplished, assess the number of supercomputers required to meet the realistic needs of the community each agency supports both for scientific research and graduate education;

2. develop plans to ensure that high quality service is provided to their research communities to meet the needs;

3. begin by expanding the user community of some of the existing supercomputer centers that are (i) already providing high quality service to remote users or (ii) have significant potential for achieving this.

4. initiate access by common carriers for communications by remote users. In order to provide initial service, 9600 baud lines should be made available to researchers.

To meet long-term objectives, each agency should:

5. coordinate the development of plans for scientific networks with other agencies. Over the next several years, the goal should be to design and implement an integrated system of networks linking researchers with each other and with computing centers.

6. establish new supercomputer centers and associated networks as needed. Some of the considerations for managing and siting centers described in this report may be of use to the various agencies. Existing centers should be used to help train personnel for new centers.

7. consider carefully the value of initiating new centers versus the economic savings of expanding existing centers. Economics and the present scarcity of trained people suggest that the number of supercomputer centers should increase only as staff are trained and needs develop.

8. provide necessary funding of supercomputer centers to maintain competent personnel and to provide high quality service. The most practical means for avoiding inefficient fluctuations in funding is for agencies to provide direct funding to meet anticipated center costs.

9. begin experimenting with high bandwidth communications systems. High bandwidth communications will play an important role in supercomputer networks in the future.

10. support graduate student training in the use of supercomputers and as appropriate in the operation of supercomputer centers.

And, finally, it is recommended that:

11. federal agencies should continuously upgrade the capacity and capabilities of federally supported supercomputer centers; and

12. a mechanism should be established to trade computer time between agencies. The objective of such trading would be to permit a researcher who has contracts with more than one agency to use only one computer system, thus saving time and effort.

REPORT OF THE FCCSET WORKING GROUP ON ACCESS TO SUPERCOMPUTERS

II. Present Status of Scientific Computing

Background - The Issue

During the decade of the 1970's, the government laboratories rapidly began to acquire supercomputers, the most advanced general purpose computer system available. These supercomputers were used as powerful tools for conducting research in fusion energy, weapon systems, aerodynamics, and atmospheric science. They were the Control Data 7600, Burroughs Illiac IV, and Texas Instrument Advanced Scientific Computer. They were noted for their greatly increased speed and memory over previous generations of computers.

The introduction of the CRAY-1 and CYBER 205 in 1977 and 1981, respectively, substantially boosted the capabilities of supercomputers, thus attracting more and more researchers to use the tools for numerical modeling and processing of very large data sets. Today, fully configured, dedicated systems with memories of at least 1 megaword and speeds of more than 100 MFLOPS exist at a number of government laboratories and research centers. Many of these systems are saturated and collectively serve only a limited portion of the nation's research community. Just two centers, National Magnetic Fusion Energy Computer Center (MFE), and the computing facility at the National Center for Atmospheric Research (NCAR) have served, in recent years, a substantial number of remote researchers.

In the Federal Government, supercomputers are used primarily for weather forecasting, space, energy, nuclear weapons, and science research projects. This is summarized below:

> Department of Energy. The primary use of supercomputers in DOE laboratories has been for modeling nuclear weapons performance. Modeling has substantially reduced the number of experimental tests required to finalize design of new weapons. Historically, the DOE weapons laboratories have been the world's largest user of supercomputers and they have been instrumental in developing software to utilize new generations of machines.
>
> In more recent years, applications of supercomputers have increased in various energy technologies including fusion research and nuclear energy.
>
> The energy sciences programs (High Energy and Nuclear Physics, Basic Energy Sciences, and Health and Environmental Research) also have large complex scientific problems and are just beginning to benefit substantially from the use of supercomputers.
>
> National Science Foundation. Many university researchers have found that in order to understand important aspects of physical, biological, and chemical events and processes, numerical simulation on supercomputers is critically needed. In recent years, NSF has provided supercomputer facilities for atmospheric and ocean sciences at NCAR. But only small numbers of university scientists and engineers in other disciplines have had access to supercomputers at other national laboritories, and only if their research was relevent to the laboratory mission. Much of the university based work in numerical modeling has been accomplished on low cost mini-computer systems.

National Aeronautics and Space Administration. Through the development and deployment of remote sensors, the U.S. space science program has been gathering enormous data sets on the earth and solar system. Advances in sensor technology, communications, and computational capabilities have been important ingredients in this program. Within the space science data management environment, supercomputers are one of the essential elements. The steady increase in data volume, variety of space-derived data, subsequent processing requirements, and modeling needs of scientific investigations have greatly increased the need for the fastest and most capable computers. NASA's research and technology pathfinding efforts in the computationally based sciences, particularly computational fluid dynamics, has put the U.S. in the leadership role in several disciplines. These forefront research efforts require access to the leading edge computational capabilities. As a result of continued improvements in numerical algorithms as well as dramatic increases in processing speed and storage capacity of recent supercomputers, computational fluid dynamics has become a powerful and widely used tool in the aerodynamic design of aircraft and aerospace systems.

Department of Commerce. Supercomputers are being used extensively in the weather forecasting and meteorological research activities of the Department's National Oceanic and Atmospheric Administration. The dynamic nature of the atmosphere and its effect on human activity has led to the development of complex technologies in weather observing and forecasting capabilities. Supercomputers are used for the processing and dissemination of meteorological data and for the developing of forecasts. To better understand the physical laws that govern the behavior of the atmospheres and oceans, sophisticated numerical models have been developed. Modeling research is heavily dependent on supercomputers to compute the complex interactive processes with the extensive space and time scales involved, and the large data input required to establish initial conditions.

Coupled with this federal activity, industry has successfully used supercomputers in a number of applications including: exploration for petroleum and mineral deposits; maximizing petroleum recovery; design and simulation of aircraft components; design and simulation of complex integrated circuits; design of future supercomputers; chemical process control and optimization; combustion modeling for internal combustion engines; and computer graphics for image generation and animation.

In spite of this steady growth in supercomputer capabilities and their availability to some research groups, two very serious problems exist and are threatening to affect the vitality of this nation's scientific and engineering research and technological base. They are:

1. Large numbers of researchers and potential researchers (graduate students, etc.) do not have effective access to supercomputers; and

2. The capabilities of today's supercomputers are several orders of magnitude too small for attacking many current and urgent research problems.

This critical situation was identified by the Panel on Large Scale Computing in Science and Engineering, chaired by Peter D. Lax. In June of 1982 the NSF and DOD (with the cooperation of DOE and NASA) sponsored a workshop for this panel to explore the problems, needs, and opportunities in large scale computing. In the December 1982 report of this panel (Lax Report) it was stated that:

> "The U.S. has been, and continues to be, the leader in supercomputer technology and in the use of supercomputers in science and engineering. The present position of leadership is evidenced by the dominance of the supercomputer market by American producers and by the successful exploitation of supercomputing at national laboratories. However, the Panel finds that this position of leadership is seriously undermined by the lack of broad scale exploitation, outside of a few national laboratories, of the scientific and engineering opportunities offered by supercomputing, and by a slowdown in the introduction of new generations of supercomputers. This threat becomes real in light of the major thrust in advanced supercomputer design that is being mounted by the Japanese Government and industry, and by vigorous governmental programs in the United Kingdom, West Germany, France, and Japan to make supercomputers available and easily accessible to their research and technological communities.[1/]

As a result of this assessment, the Lax panel recommended the establishment of a National Program on Large Scale Computing. This program should foster participation and cooperation by federal agencies, universities, and industry in carrying out four goals:

- increase access for the research community to supercomputers;
- increase research in the computational tools (software, algorithms, etc.) needed for effective use of supercomputers;
- train personnel in use of large scale computing; and
- stimulate R&D for substantially increased capability and capacity in the next several generations of supercomputers.

[1/] See page 6 of the Report of the Panel on Large Scale Computing in Science and Engineering, (December 1982).

In addition to the Lax report, there are a number of other relevant studies some of which are listed in Appendix B.

Several of the European nations have moved aggressively to supply researchers with access to supercomputers. It is not uncommon for U.S. researchers who have difficulty obtaining supercomputer time in their own country to do supercomputing on European machines during summer visits, sabbaticals, etc.

A list of countries and machines entirely or partially available for university researchers is provided below. This list may not be complete.

Country	Site	Computer
Germany	Max Planck	CRAY-1
	Bochum	CYBER-205
	University of Karlsruhe	CYBER-205
	University of Stuttgart	CRAY-1
	KFA	CRAY-1
France	Ecole Polytechnique	CRAY-1
England	Daresbury	CRAY-1
	University of London	CRAY-1
	University of Manchester	CYBER-205
Sweden	Volvo	CRAY-1 (1/2 for academia)
Japan	University of Tokyo	HITACHI S220

In Japan, supercomputer companies are placing their first production machines in university centers.

In December 1982 the Science Advisor, Dr. Keyworth, initiated a review of U.S. policy in the field of supercomputers by the Federal Coordinating Council on Science Engineering and Technology, FCCSET. As part of the review, this interagency working group was established to consider the question of research access to supercomputers and to make recommendations to the FCCSET. Two parallel working groups were established to review the other aspects of the general topic. This report examines the access issue.

During its deliberations, the Access Group heard presentations on the operation of the Magnetic Fusion Energy Computer Network and the National Center for Atmospheric Research's Computer Network. These two networks provide supercomputer access to the largest numbers of remote users of any existing supercomputer centers. The Group also heard agency plans for providing increased access to supercomputers in the future. These presentations included: a report by NSF working group on plans to provide significantly enhanced opportunities for the academic research community to access supercomputers; NASA's plans to provide more access to its contractors as part of the Numerical Aerodynamic Simulation program; DOD's plans for the ARPAnet and the wideband packet satellite system; and DOE's plans to expand access to the Magnetic Fusion Energy (MFE) net for contractors funded by the Department's Energy Research programs.

Two members of the group, L. Lee and J. Decker, together with K. Curtis and A. Thaler of NSF, and S. Karin of GA Technologies (serving as a technical consultant) visited supercomputer centers at the University of Minnesota, NCAR, NASA-Ames, and the Lawrence Livermore National Laboratory (MFE). Discussions were also held with a number of scientists who are remote users of the NCAR center.

III. Capability and Service Requirements of Supercomputer Centers

High Quality Service

It has been well documented that supercomputers are critical in attacking many difficult scientific and technological problems. The remote use of supercomputers should be made convenient and friendly to the user. Otherwise the inefficiencies at the user end can be substantial, making the real cost of computing very high. In fact, researchers may not use this important scientific tool if its use proves difficult to learn, frustrating and generally inconvenient. The Access Group heard anecdotal evidence of poor supercomputer service causing researchers to return to work on their local minicomputers. For these reasons, the Group wishes to stress the importance of providing high quality service.

As mentioned earlier, several members of this FCCSET access working group visited the University of Minnesota, NCAR, NASA-Ames, and MFE. The purpose of the trip was to gain an understanding of the factors that make for successful supercomputer centers. On this trip, there were several opportunities to discuss those factors with users. The group concluded that the following features are considered essential for providing users with high quality supercomputer resources.

- medium speed (4800-9600 bps) terminal access and high speed (56 kbps or greater) file transfer capability for access by remote users;
- trillion bit or greater on-line mass storage capacity;
- full-time user consultants on software, hardware, and communications;
- extensive system documentation and software utilities and libraries;
- state-of-the-art graphics tools;
- around the clock operations; and
- interactive operating systems.

Computer Networks

A high quality, high bandwidth computer network is one of the most important requirements for providing effective supercomputer service to remote users. Networks provide an important service that goes far beyond the simple concept of providing supercomputer access to remote users. Geographically dispersed researchers can collaborate daily using electronic mail and shared program and data files to work together effectively on the same problem. Duplication of programming effort can be virtually eliminated if collaborators can easily share resources. The cost of the communications network in use on the NMFECC is a small fraction of the total operating budget, yet in terms of the valuable time spent by a scientist using this facility, high bandwidth response is a crucial part of the service. The 1200 bps terminal access service commonly available today is not adequate for the large majority of supercomputer users. A minimum of 9600 bps is necessary to provide the level of response available on local minicomputers; there is no reason why this level of response should not be available on supercomputers.

For file transfers of data and programs between local computers and supercomputer centers, a minimum of 56 kbps is usable, but 1.5 Mbps is both desirable and feasible. Because the efficiency of the researchers using the centers is by far the most important consideration, and because the cost of communications is a small fraction of the total cost of providing these services, the highest quality network available should be obtained.

Special Characteristics of Supercomputer Centers

Supercomputer centers are defined here to mean centers equipped with one or more Class VI computers; an adequate hierarchical mass storage system; state-of-the-art output facilities; a full complement of systems, operations and user services staff; and network support for remote users throughout the U.S. Two such centers, NMFECC and NCAR, are supported by government research programs to provide high quality service to a large number of researchers in universities, industry, and laboratories. In both cases, it required several years to build up competent staff; to assemble complete supercomputer systems, and to develop operating systems and communications software. In the case of the NMFECC, a significant portion of the original staff were recruited from the LLNL central computer center, a classified center serving a large number of internal users. Even with the experienced initial staff, it took several years to develop good service to remote users.

In terms of skilled staff, system hardware and software, and communications facilities, the requirements for the operation of a successful center serving a large number of remote users are quite different from those for a center operating in a batch mode for on-site and dial-up users. Although it is certainly possible to expand some of the small supercomputer centers to provide high quality service to a large number of remote users, it will require a significant amount of time for system development and hiring and training of competent staff.

Some effort on software development could be saved by duplicating parts of an established system such as the NMFECC. However, some of the components used in these centers are now no longer manufactured and new hardware and software development will be required to replace these components.

Supercomputer vendors have not provided adequate software in the past. DOE laboratories, the largest users of supercomputers, have developed their own operating systems to provide a high quality, interactive computing environment. New supercomputer centers acquiring new classes of supercomputers may either have to develop the capability of writing an operating system or wait until another center has developed one. This will be required unless the vendors commit resources to develop satisfactory interactive operating systems.

As indicated above, operation of a supercomputer center to provide high quality service to remote users is not a simple step up from operation of a general purpose mainframe that comes with good software support and a complete set of peripherals from large commercial vendors. In order to reduce the learning time for any new supercomputer center, there should be a mechanism for the transfer of knowledge from existing centers. One possibility is to send staff from a new center to existing centers for training.

The cost of operating good supercomputer center service for remote users is reasonably well established. Although costs will depend on the exact nature of the service provided and with the institution providing the service, both NCAR and NMFECC can be used as indications of costs. The annual operating budget of the NCAR center is approximately $10,000,000. This center operates two Cray-1's with a staff of 100 and serves about 1000 remote users. The total value of installed equipment is about $30,000,000.

The total annual operating cost of the NMFECC is approximately $12,000,000. This center operates two Cray-1's and a CDC 7600, and serves about 2000 users. The total value of the installed equipment is $25-30M. An approximate breakdown of the NMFECC budget for FY 1983 follows:

Manpower (80 FTEs)	$ 6,400,000
Lease (1 Cray)	2,700,000
Maintenance	1,300,000
Communications	800,000
Supplies and Service	900,000
Miscellaneous	400,000
Total	$12,500,000

IV. Some Conclusions on Operating and Siting of Supercomputer Centers

The list of siting considerations provided below is not to be taken as all inclusive, nor is it intended to be a strict set of criteria that agencies should follow. The list does represent some observations and thoughts of the Working Group which should be carefully considered. Many of the considerations come from examining existing successful operations.

(1) The proper philosophy of operation of a supercomputer center is absolutely essential for providing good services to a large collection of remote users. Such a center has to give the highest priority to satisfying user needs. In particular, the desire to experiment with the system in a research mode is incompatible with serving users. The service mission must be compatible with the goals of the institution that manages the center.

(2) The center must receive frequent, organized input from users. Both NCAR and MFE have advisory committees which are credited with contributing substantially to the success of those centers.

(3) Stable funding is an important ingredient of success. Computer centers have many fixed operating costs such as equipment leases and maintenance, staff salaries, communications, utilities, supplies, etc. There are two primary options for meeting these costs: direct charges to users or direct funding of the center. In the former case, computer time can be purchased by researchers who have the money and choose to do so. In the latter case, computer time is allotted by the funding agency or a peer review committee acting on behalf of the agency. Direct charges to users does not guarantee stable funding. In fact, if projected revenues are not being met, the center has to increase the hourly rate. This may drive customers away, further reducing revenues. Note that a center serving remote users is in a more precarious position than a center only serving users within an institution. In the latter case, the management of the institution can obtain a commitment from researchers to spend funds at the local center and avoid a disasterous financial situation. The two centers with the longest record of success in serving remote users, NCAR and MFE, both receive direct funding.

In the case where an agency funds a computer center directly at the beginning of the fiscal year, computer time is allocated by the agency in a manner similar to program funds. When a principal investigator submits a research proposal, he may request resources consisting of supercomputer time as well as operating funds. Requests for computer time are reviewed in the same manner as operating funds. The normal situation is that requests for computer time like those for program funds far exceed the totals available and consequently both have to be managed carefully as scarce resources.

(4) There is merit in locating one or more supercomputer centers near or in universities to serve educational needs. However, since supercomputer centers are costly to establish and operate, it is clear that only a very few centers are possible. It thus seems appropriate that university consortia might be formed to operate centers in several regions of the country. Further, the consortia could establish or hire management organizations who have as their primary objective providing computer service. The two centers with the longest record of success are both run by university contractors but in an indirect fashion. MFE is run by the University of California through the Lawrence Livermore National Laboratory and NCAR is managed and operated by the University Corporation for Atmospheric Research--a consortium of 49 universities.

(5) There is a considerable cost saving realized by locating several supercomputers at one center compared with establishing several small centers. The incremental cost of adding one more computer to a center is generally small compared to establishing a new center. One center can economically serve a large geographical area since communications costs are a relatively small fraction of the total operating cost. For the MFE net, the communications costs are a small fraction of the total. Furthermore, supercomputer centers once established should be maintained at the state-of-the-art, which means replacing mainframes every five or six years. It is questionable whether funds will be available in the future to keep a large number of centers up to date.

While the economics of small supercomputer centers, e.g. one supercomputer, are less favorable than for larger centers, there is also an upper limit to computer center size. It is not clear what determines that limit. Some possible factors are: number of computers, size of staff, number of users, number of disciplines served, number of government departmental programs served, etc. To date, successful Government funded supercomputer centers have served single mission programs and/or single disciplines.

(6) Location of supercomputer centers in a good science and engineering environment where computing is done routinely is a definite plus. Such nearby activities help computer center staff maintain direct contact with the problems of their users.

(7) Since supercomputers represent leading-edge technology, a supercomputer center must have the capability to test and implement new software and hardware concepts.

(8) A center must provide good user documentation, training, and full-time user assistance.

(9) A center must be able to attract high quality personnel to carry out many of the items listed above.

V. Goals for U.S. Policy

Supercomputers are very powerful scientific tools that have applications in many areas of research and in some fields have become essential for advancing the state-of-the-art. In order that the U.S. retain its technological leadership in these fields, supercomputers must be made available to the many researchers in university and government laboratories across the country who currently do not have access. Toward this end, the goals of U.S. Government policy should be:

(1) To provide access to supercomputers for researchers who are conducting research supported by or of interest to the government and who require such capability; and

(2) To aid the education of university students in the use of supercomputers.

The implementation of a policy based on this goal will be instrumental in:

(a) Ensuring that the U.S. retains its lead in many areas of science and technology by providing researchers in the country with the use of a very important tool;

(b) Ensuring that the U.S. maintains its lead in the development of advanced supercomputer systems by expanding the applications, and in turn, the market for supercomputers in the U.S.; and

(c) Providing the opportunities for training scientists and engineers in the use of supercomputers to meet future needs for experienced researchers in Government programs and U.S. industry.

VI. Recommendations

The Government should take action to provide access to supercomputers for government funded researchers in government laboratories, universities, and other institutions who have large computational problems requiring the use of supercomputers. Supercomputers are critical in solving the complex problems vital to our security, economic and scientific interest. There is clear evidence that the lack of access to supercomputers limits progress in many aspects of science and engineering.

The Working Groups's recommendations were developed to help solve this national problem. Recommendations are made on near-term approaches for providing access to large numbers of users in the shortest time and also for the longer term. The Group tried to identify those features of the MFE and NCAR operations (described in Appendix A) which have made them successful in providing high quality services to remote users. Recognition of those features was a key factor in developing many of the recommendations.

Recommendations to meet Immediate Needs

A continuing interagency group should be designated to provide coordination among agencies to help ensure appropriate sharing of resources, to avoid duplication of effort, and to maintain a dialog on policy and technical issues.

1. The number of supercomputers required to meet the needs of each agency should be assessed if such an assessment has not already been conducted.

2. Each federal agency should develop plans to ensure that high quality service is provided to its own research community. High quality service requires good communications and well-managed supercomputer centers dedicated to service.

3. In order to accomplish the above recommendation, some of the existing supercomputer centers already providing quality service to remote users should be used to provide service to a wider community in the near term. For example, The Department of Energy is taking this approach by expanding the MFE network to serve the energy sciences programs.

4. In order to initiate access as soon as possible, existing communications should be used as a start, including common carriers, ARPAnet, MFEnet, etc. In order to provide initial service, 9600 baud lines should be made available to researchers.

Recommendations to meet Long-Term Objectives

5. Federal agencies should coordinate the development of plans for scientific networks. Over the next several years, the goal should be to design and implement an integrated system of networks linking researchers with each other and with computing centers.

6. Establishment of new supercomputer centers and associated networks should be initiated by individual agencies as needed to meet their long-term objectives. The benefits derived from adding supercomputer centers include opportunities to improve education and training and to optimize the coupling of researchers to supercomputers. Some of the considerations for siting centers described in this report may be of use to the various agencies. It must be recognized that it will likely take several years before a new center is providing high quality service. The use of existing centers to train personnel for new centers would aid development of new centers substantially.

7. For a center with two supercomputers and a total installed equipment base of $25-30M, annual operating costs are likely to be $10-$15M per year. Because the total costs of starting a new center are higher than the incremental costs of adding a machine to an existing center, the total number of centers should be limited and should increase only as staff are trained and needs develop.

8. Stable funding of supercomputer centers is an essential ingredient in maintaining competent personnel and for providing high-quality service. The most practical means for realizing truly stable funding is by agencies providing funding to meet center costs. The agencies would then allocate computer time in the same manner as they allocate program funds. This approach has been used successfully by NSF (NCAR) and DOE (MFE). Direct charging of users particularly when they are remote will not provide the necessary stability of revenues.

9. High bandwidth communications will undoubtedly play an important role in supercomputer networks in the future for applications such as interactive graphics. The federal agencies should begin experimenting with high bandwidths communications systems.

10. Each agency should support graduate student training in the use of supercomputers. Fellowship programs and summer colloquia should be established. Special allocations of supercomputer time should be made available for projects that involve graduate students. Supercomputer centers should be encouraged to employ students on a part-time or full-time basis wherever appropriate to provide practical experience in the management, operation, and systems software demands of providing supercomputer services.

Additional Recommendations

11. Federal agencies should continuously update the capacity and capabilities of federally supported supercomputer centers. Supercomputers are by definition the most powerful scientific computers available at any time. This implies that a supercomputer center must have the latest state-of-the-art machine to meet the evolving computational requirements of the research community.

12. A mechanism should be established to trade computer time between agencies. The objective of such trading would be to permit a researcher who has contracts with more than one agency to use only one computer system, thus saving time and effort.

APPENDIX A

Current Status of Scientific Access to Supercomputers

The Lax Report emphasized that increasing access to supercomputers requires high bandwidth networks, fully-configured new and updated supercomputer centers, and opportunities for experimenting with new computers. Indeed, all of these elements are essential. To understand the extent of which these elements exist today, the following summarizes the key characteristics of networks and supercomputer facilities that are now operational and which serve remote users.

Computer Networks

In the late 1960's the Defense Advanced Research Projects Agency (DARPA) began an international packet switched computer network called the ARPAnet. Today that network hosts over 200 computers at nearly 100 universities, government laboratories, and private research firms. The backbone of the network is 50,000 bits per second land line augmented by satellite links to Hawaii and Europe. ARPAnet has been split into an operational network, called MILNET, and an experimental network, called ARPANET.

The ARPANET is the only operational network that provides the full range of network services (virtual terminal access, file transfer, electronic mail and bulletin boards, and internetworking protocols) for a large variety of host computers. This serves as an example of the kind of services desirable for a future operational network for supercomputers.

A second example is a wideband packet satellite network which was developed with DARPA support to explore issues in high-speed computer communications and real-time data transmission including integrated data and voice. The network operates across the continental U.S. using a single broadcast channel on the Western Union WESTAR satellite. The system operates coded or uncoded at data rates up to 3.088 M bits/sec.

A third example is the satellite network run by the NMFECC. This provides dual 56kbps channels from 4.8 kbps and 9.6 kbps access lines into the regional centers. The direct satellite link provides reliable, medium band width file transfer capability as well as terminal access.

Commercial value added networks provide only low-speed terminal access (up to 1200 bps). The packet switched networks, TELENET and UNINET, evolved from the ARPANET. TYMNET, a circuit switched network, evolved independently, but also provides only terminal access. Other networks, such as NSF's CSNET, Bell Labs' USENET, and EDUCOM's EDUNET, are built on the ARPANET, the commercial networks, and dial-up service. These serve various communities and terminal access and electronic mail.

Computer networks provide an important service that goes far beyond the simple function of allowing remote users terminal access to supercomputers. Networks allow geographically dispersed researchers to collaborate effectively by sharing software, trading ideas, and even jointly writing research papers.

National Magnetic Fusion Energy Computing Center (MFE)

For almost ten years now, the MFE computer center has been providing fusion energy researchers with large scale computing services. Scientists are developing several categories of computer simulations to study the physics of magnetically confined plasmas, while engineers are designing the models needed in fusion reactor design studies. The MFE provides high level supercomputer capability to about 2000 users at more than 30 remote locations. At the heart of the center is a CDC 7600 plus two CRAY supercomputers with one and two million words of memory, respectively. The file management system has three levels of service - disk systems, mass store tape cartridges, and automatic tape library.

The cornerstone of the MFE system is the network. Four major DOE fusion energy laboratories are linked by dual 56 kbps satellite links to MFE. Many universities and other research laboratories are connected to these centers by medium speed land lines (4800 - 9600 bps). Other users, not linked by one of the above, may dial up the center through TYMNET, ARPANET, direct commercial dial, or FTS. This nationwide network is efficient, cost-effective, and allows remote fusion researchers to access the center with several levels of local computing capability.

There are numerous aspects of the MFE center which contribute to its success in serving fusion energy researchers. Some of these are enumerated below:

- Growing national network consisting of a hierarchy of communication speed and hardware capabilities;

- Fully interactive time-sharing system on the CRAY computers;

- Consultants to users who specialize in mathematics, physics, and engineering software and who have experience in hardware; and

- Capability of developing new services such as an interactive multiprocessor operating system for supercomputers, new graphics tools, and an advanced interactive query system.

National Center for Atmospheric Research (NCAR)

The National Center for Atmospheric Research has been providing computational services to the atmospheric science community for twenty years now. The acquisition of successive generations of computers has made it possible for the science to address problems in the theory of weather prediction, climate studies, severe storms research, ocean circulations, and the influences of the sun, land, and oceans on the atmosphere. Today, NCAR provides computational resources to more than 900 scientists at 75 institutions. The hardware at NCAR includes two CRAY computers (1 megaword memory, 64 bit word lengths), two IBM 4341 interactive computers, and a mass storage system on which are stored 7.6 trillion bits of data. In addition to NCAR's important role in serving the computational needs of the atmospheric science community, NCAR has become a leading computational resource for the ocean science community.

One of the most significant strengths in the management and operation of the NCAR computing facility is the involvement of the user community. NCAR is managed by the University Corporation for Atmospheric Research (UCAR) - a consortium of 49 universities. The atmospheric and oceanic research communities (NCAR and university scientists) participate in the allocation of computing resources, setting of plans and priorities, and reviewing of policies and operating procedures. This has led to the development of the following very extensive and useful services:

- On-line access to many large meteorological data sets
- Large consultant group and educational program
- Multiple capabilities in graphical displays
- Numerous numerical and utility software libraries

There is a great need for a cost effective, efficient, and low-error rate communication system. Much of the computing on the NCAR system is data intensive. NCAR has recently taken some steps to improve access for its remote users. As a result of an NCAR agreement with the national communications network UNINET, 95% of the researchers at universities can communicate with the NCAR system at 1200 baud and at the cost of a local telephone call. Plans in the future include the development of a high bandwidth communication system linking university researchers to NCAR.

Los Alamos National Laboratory

One of the largest scientific computer centers is at Los Alamos National Laboratory. Currently, the laboratory serves 4000 users in Los Alamos and 1000 remote researchers. The primary function of the computer facility is to serve the computational needs of the nuclear weapons program. However, Los Alamos also supports the computational needs of researchers in nuclear and plasma physics; inorganic, nuclear, and physical chemistry; mathematics and computer science; metallurgy; biomedical effects of radiation; and earth and space sciences. The nucleus of the computing facility is a collection of interactive computers and high capacity, rapid access storage devices. These include:

- 5 CRAYs
- 4 CDC 7600s
- Numerous CYBER 800's and VAX 11/780s
- IBM 3850 Mass Storage System

This network of computers and data storage devices is partitioned into three security partitions: secure, administrative, and open.

The Los Alamos Computer Center performs research, development, and support in areas of computer science and engineering. A very high-speed interactive display system of both graphical and textural data is being developed. Also, the center supports five operating systems, all of which provide the user with interactive computing. Recently a Denelcor HEP-1 multiprocessor computer was acquired for experimentation. The center is studying several high-level languages to determine their usefulness in solving a variety of scientific problems. In addition to providing users with many direct services, such as consulting, troubleshooting, assisting with equipment, and writing, maintaining, and enhancing a variety of libraries, the center is developing a users' education program centered around classes, seminars, computer-assisted instruction, and videotapes.

The Computing Division at Los Alamos recently began to provide computing support to over 400 external users associated with the Defense Nuclear Agency (DNA). It also supports some supercomputer access from university research groups. Remote access to Los Alamos is supported through commercial dial-up lines, TELENET, ARPANET, and Los Alamos' own extended network (XNET). VAX computers with the DEC VMS operating system are required for using XNET. Also, six 56 kbps lines have been implemented to serve remote DNA users.

NASA-Ames Research Center

NASA is planning to develop, at its Ames Research Center, a Numerical Aerodynamic Simulation (NAS) system. This is a unique supercomputer facility for solution of aerodynamic and fluid dynamic problems. In 1985, the NAS will become operational and provide these supercomputer system services to NASA, DOD, industry, other government agencies, and universities.

In mid-1982 NASA significantly changed its approach for the development of the NAS. The former approach was for NASA to procure a contract for the development of a very high performance supercomputer system specifically designed for solving viscous, three-dimensional, fluid-flow equations. In the latter phases of this development, neither of the two competitive vendors could assure NASA that they could meet their requirements on performance. Under the new NAS approach, Ames will acquire in FY 1984 a prototype of a Class VII large-scale computer. This should be a 250 megaflop, 64 megaword main memory system. This would be followed by a second procurement in FY 1987 of a 1 gigaflop (1000 megaflops) and 256 megaword Class VII computer. The NAS program calls for a systematic and evolutionary approach to acquiring computer-system technologies to maintain a leading-edge performance capability. Hence, for example, users will be accessing one supercomputer system while NASA works on procuring and installing a significantly better performing follow-on system.

There is a need for NASA to provide the full end-to-end capabilities. This will span the range from mass storage systems to microprocessor work stations, and will provide landline and satellite data communication links to offsite users.

The overall cost of the NAS in the FY 1983-88 time frame will still be in the neighborhood of $100M. However, in this new incremental approach to developing the NAS, users will have access much earlier, FY 1984 versus FY 1988. One interesting and appealing aspect of this plan is that there is a good possibility that as the NAS evolves, it will become a multi-vendor complex of supercomputer systems. This may prove to be very advantageous to a diverse aerodynamic industry user base.

NASA-Ames is currently operating a CRAY-1S and a CDC 7600 for the internal needs of the research center. This fall the center will replace the CRAY-1S with a CRAY-XMP. A study is being initiated on the needs of long haul communications for remote user groups of the NAS.

Other National Laboratories

FACILITY	MAJOR SYSTEM	RESEARCH USERS
NASA-Goddard	CYBER 205	NASA-funded and NASA-project related
NASA-Langley	CYBER 203	NASA and NASA-funded scientists
NASA-Lewis	CRAY 1-S	Principally Aerodynamics-related
Geophysical Fluid Dynamics Lab	CYBER 205	Laboratory Personnel
Air Force Weapons Lab	CRAY-1	Local and Remote
Naval Fleet Numeric Center	CYBER 205	Weather Forecasting and related Research

Universities

Several Class VI computers have recently been acquired at major universities - a CYBER 205 at Colorado State University (CSU) and Purdue University and a CRAY-1 at the University of Minnesota (UM). At least two of these systems are seriously underutilized. After more than one year of availability the Colorado State system, it is being used at approximately 30% of its capacity. The University of Minnesota has upgraded their CRAY from 1/2 million words to one million. All three universities have received some financial assistance from vendors. These university centers are striving to increase staff experience in serving large numbers of remote researchers and to develop their communications systems for remote users.

Commercial Services

Supercomputing services are also available through commercial service bureaus. Boeing Computing Services and United Information Services provide time on CRAY-1 supercomputers and Control Data Corporation on a CYBER 205.

APPENDIX B

Additional Supercomputer Studies

- The report on Computational Facilities for Theoretical Physics (June 1981) recommended that a national computational physics network be established with a supercomputer facility as a principal node and a set of regional and local nodes, each with its own computing capacity.

- The 1982 NAS report entitled An Assessment of Computational Resources Required for Ocean Circulation Modeling surveyed the ocean modeling community and found that the projected computing resources of their academic community will fall short of their needs in 1984 by a factor of more than two. As a result of this report, NCAR has been selected to take the responsibility for NSF supported oceanographic computing.

- The Atmospheric Sciences Advanced Vector Computer Justification document (December 1982) stated a need for increased computational resources at NCAR by a factor of 5 over their presently available two Cray's.

- The computational fluid dynamics community has concluded in the NRC Report on Computational Fluid Dynamics (1983) that they need a 1000 MFLOP computer with a 500 million word memory in order to simulate the flow field about a complete aircraft using the Reynolds-averaged Navier Stokes equations.

- In the NRC Report on Data Management and Computation (1982) there are a number of recommendations to NASA on the management of space science data. Among them, it is recommended that NASA work closely with the scientific community to assure access to adequate computational capabilities, communication facilities, and software.

- In January 1983 the DOE report Magnetic Fusion Energy and Computers, a panel of researchers recommended that the MFE acquire in the next several years additional supercomputer systems (Class VI and a Class VII) and that the communication systems be upgraded by adding more dedicated lines and increasing bandwidths.

- The NSF report A National Computing Environment for Academic Research recommends that up to 10 new supercomputer centers be established to serve the documented needs of the research community.

FCCSET Access Group

James F Decker
James F. Decker
Department of Energy

Robert E Kahn
Robert E. Kahn
Department of Defense

Larry Lee
National Science Foundation

Richard Dunlavey
NBS/ICST
Department of Defense

Ken Wallgren
National Aeronautics and
Space Administration

George Rogers
George Rogers
Intelligence Community Staff

Report to
The Federal Coordinating Council
on Science, Engineering and Technology
Supercomputer Panel
on
Recommended Government Actions to
Retain U.S. Leadership in Supercomputers

I. Executive Summary

A. Background

In January 1983, a Federal Coordinating Committee on Science, Engineering and Technology (FCCSET) Panel on Supercomputers was formed to examine what, if anything, the U.S. Government should do to stimulate the use and development of supercomputers. As a result of the FCCSET panel discussions, three interagency working groups were formed. These groups are to examine and make recommendations on the following supercomputer issues:

1) What should the Government do to ensure that the U.S. retains its lead in supercomputers? This group is named the Procurement Group.

2) What should the Government do to make supercomputers available to more researchers, particularly in universities? This group is named the Access Group.

3) A third group was formed to provide coordination among Government funding agencies research contributing to the technology base.

This Report describes the findings of the Procurement Group. The Procurement Group consists of the following members:

James F. Decker, DOE (Chairman)
George Deskin, DCI
Leonard A. Harris, NASA
E. F. Infante, NSF
Robert E. Kahn, DOD
Joe Smagorinsky, DOC

This Group was assisted by a DOE/NASA working group with the following membership:

James Decker, (DOE)
Dave Nelson, (DOE)
Bob Greeves, (DOE)
Don Austin, (DOE)
Glenn Kuswa, (DOE)
Bill Bostwick, (DOE)
Kate Foley, (DOE)
Paul Schneck, (NASA) -

Before developing their recommendations, these committees discussed many issues of supercomputer development, use, and markets with computer scientists from Government laboratories and universities. Discussions were also held with senior management of the supercomputer vendors.

Supercomputers are the most powerful general purpose computers available for large scale scientific computation. At the present time, the Cray-1, Cray-XMP, and the Cyber 205 are the most advanced machines available. Fifth Generation computers are computers based on artificial intelligence techniques and are not conventionally defined as supercomputers, even though they may require very fast processors and large memories. The development of Fifth Generation computers as defined by the Japanese was not considered by this group. Similarly, special purpose scientific computers, which perform very well on one or two specific problems, were not considered.

B. Findings

1. Supercomputers have become an essential scientific tool in many areas of Government research and development. The U.S. Government is the largest user of supercomputers, having purchased or leased over 50% of all the present class of supercomputers, Class VI machines, currently in service in the U.S. These machines are being used to model many complex physical phenomena found in nuclear weapons design, weather prediction, aerospace, magnetic and inertial confinement fusion, fundamental physics, and many other fields of research and development. Government use of supercomputers is expected to grow in the future as our ability to model complex phenomena improves. This improved modeling capability will result in continued replacement of expensive experimental testing.

2. The use of supercomputers in the private sector is increasing rapidly. Key industries such as oil, automobiles, electronics, chemicals, and aerospace are applying the supercomputer to technical problems that have important commercial implications. Future development of competitively priced commercial products with superior performance in these and many other industries will depend heavily on the availability of continually improved machines within a few years. The Japanese apparently have recognized this potential and have initiated a major supercomputer effort called the National Superspeed Computer Project. In addition to the potential international market, the Japanese are planning for a large domestic market where according to published reports, many of the supercomputers that they plan to produce will be sold for use in Japanese industry.

3. The U.S. Government and its laboratories historically have played an important role in causing each new supercomputer system to be developed. The Government has always been a "friendly" buyer by accepting certain risks associated with buying the first few machines of each new class. Each new machine has arrived at a Government laboratory with minimal operating software and devoid of any applications software. The laboratories, particularly Los Alamos and Livermore, have developed substantial amounts of software required to make these machines into useful systems.

4. Current Government programs require supercomputers with capabilities substantially in excess of those available today. Within this decade, programs in nuclear weapons design, aerospace, weather prediction, fusion research, and many areas of fundamental research and development require computers with capabilities at least 200 times greater than present Class VI machines, and requirements for even greater capability are projected for the next decade.

5. Because of the established relationship of supercomputers to leadership in both Government and private sector research and development, the U.S. can not afford to relinquish its lead in this important area of technology. Programs crucial for the defense and economic well-being of the country can not become dependent on foreign vendors and on the policies of their governments.

6. The U.S. supercomputer industry is small. Only two companies, Cray Research and Control Data Corporation, have produced the latest class of supercomputers. A third company, Denelcor, has entered the field. The industry is small because the development costs are large, the risks are high, and the market is currently limited to the most sophisticated applications. Each of the vendors are primarily manufacturers of computer mainframes, i.e., they produce little software and they do not manufacture components or peripherals (however, Cray has recently invested in their own integrated circuit facilities, and Control Data does make high performance disk drives). In the future, it is questionable whether the U.S. vendors can outperform foreign government supported national efforts by large integrated companies such as Hitachi, Fujitsu, and Nippon Electric Company, (see Appendix A). Nippon Electric has made the highest performance claims for a machine capable of 1,300 million floating point operations per second (MFLOPS), scheduled for delivery in 1985. The Cray-1 performs at peak speeds of about 100 MFLOPS and the Cyber 205 has possible peak speeds of 600-800 MFLOPS. The Cray-2, scheduled for delivery in early 1985, is projected to perform at a peak of 800 MFLOPS. Performance of these new machines on actual laboratory job mixes is, of course, an open question.

C. Recommendations

The U.S. should not relinquish its lead in this important area of technology. However, a "business as usual" approach on the part of both the vendors and the Government has a low probability of competing successfully with foreign joint government/industrial collaborations. During the past decade the capability of U.S. supercomputers has increased by a factor of only 4-5 every 4-5 years. The Japanese goal is a factor of 100 improvement by 1990. The U.S. Government should continue to play its traditional role in relation to the U.S. supercomputer vendors - i.e., the Government should continue to be a friendly buyer, develop software as necessary, and support basic computational sciences and base technology development. However, all activities associated with that role should be accelerated in a coherent manner in order to meet the challenge of foreign competition. Specific recommendations include the following:

1) The Government should set as a national goal the development of supercomputers with at least 200 times the capability of Class VI machines in this decade. The Government should provide an incentive by guaranteeing to buy at least three of each U.S. developed and manufactured supercomputer system that meets the goal. (Further work needs to be done to define more completely the 200x goal and the terms of a guaranteed buy). If successful, this process should be repeated for machines beyond the 200x class.

2) The Government should accelerate its purchases of supercomputers in order to ensure the health of the industry. Profits from sales of the present class of supercomputers will contribute to the development of the next class. The Government should continue to be a friendly buyer, particularly for prototype and production machines that represent important steps toward achieving the 200x goal.

3) It currently appears that the only feasible technical approach to meeting the 200x goal in the near term is to utilize parallel processors. Development of the software to obtain the theoretical benefits of using many processors in parallel to solve a single problem will be a significant challenge. Immediate experience with parallel processors is required. Government laboratories, and perhaps with government support a few universities, should purchase or lease several machines with interesting parallel architectures for experimentation. These "experimental" machines should be put on networks so that researchers from many institutions, including universities, can contribute to development of new languages, algorithms, software tools, and applications systems. A difficult but worthwhile goal is the development of a common operating system for the next class of U.S. supercomputers.

4) The development of high performance peripherals required for future generations of supercomputers is a well recognized problem. One of the primary difficulties is the small market that makes development of the desired hardware by normal private sector financing unattractive. Two possible solutions are a guaranteed government buy and/or government funding of research and development. A similar problem exists for high speed components. The vendors may not be able to obtain advanced high speed memory and logic circuits from U.S. suppliers in the future which could have serious implications for development of new machines. Continued government support of base IC technologies including promising new approaches to fast logic and memory circuits is important. The special requirements of the supercomputer industry should be examined in-depth.

5) Government support of long range research and development for scientific supercomputing should be increased. This research is required to develop the new approaches in hardware and software necessary for the substantial increases in computing power required in the next decade. Since most of this research would be supported in universities, an important product of this research will be the trained supercomputer architects and developers of the future.

6) The Government should take action to greatly increase access to supercomputers for researchers, particularly in universities. Improved access will not only bring this modern scientific tool to universities and improve their competitive position in research, but it will also help train new scientists and engineers in the use of supercomputers.

7) Research and development tax incentives for the industry should be explored by a Government committee with members from OMB, Department of Commerce, and Department of Justice.

8) U.S. export control of supercomputers has a significant effect on U.S. supercomputer vendors since Western Europe represents 40% of the market. Acceleration of the export licensing process should be explored by DOD and Commerce along with other appropriate agencies so that U.S. vendors are not unduly penalized in competing for foreign customers. Of course, export control of supercomputers must remain consistent with national security considerations.

9) A permanent interagency group should be established to coordinate individual agency supercomputer activities as necessary to implement the above recommendations. This group should function at both the policy and technical level.

II. Historical Perspectives

The federal government has exerted a powerful influence on the development of large scale scientific computers (supercomputers). The first electronic computer, the ENIAC, was built by the University of Pennsylvania under contract to the U.S. Army Ballistics Research Laboratory, and the first program run on the ENIAC was a nuclear weapons design calculation by the Los Alamos Scientific Laboratory. In fact, the Manhattan Project and the ENIAC project were initiated the same year (1943). In the succeeding four decades, the defense and computer communities have had a mutually beneficial relationship: defense programs have provided much of the motivation and funding for the development of supercomputers, and the computer industry has provided the tools for solving the highly complex problems that would otherwise have remained intractable.

Fields that have been and continue to be strongly dependent on computational modeling include nuclear weapons design, magnetic fusion energy, cryptographic analysis, aerodynamics, integrated circuit design, inertial confinement fusion, reactor safety, atmospheric research and weather forecasting, astrophysics, molecular biology and chemistry, and fundamental physics research.

The speed of computation that can be achieved in the fastest computers available at any given time has increased by some seven orders of magnitude from the early 1940s to the present. Three of these orders of magnitude resulted from the transition from electromechanical to electronic circuits. Another three orders of magnitude have been achieved in the speed of the electronic components (for example, the IBM 701 of the early 1950s had a cycle time of 12 microseconds, and the Cray-1 has a cycle time of 12 nanoseconds). The other order of magnitude in this overall increase has come about through increasing levels of parallel execution of computational tasks. Equally dramatic advances have come about through improvements to the algorithms used to solve complex problems.

The supercomputer industry is now facing a demand for several orders of magnitude increase in performance in this decade, but the projected performance increases available from faster components appear to be limited to at most one order of magnitude. Hence, it is clear that new and explicit forms of parallel computer organization will be required to meet the needs of the nation. Achieving success in this radical change in the design of supercomputers will require the cooperation of some of the best minds in the nation.

The early years of the history of supercomputers were characterized by a strong interaction among government, universities, and industry that guided and accelerated the development of very powerful scientific computers. The mathematician John von Neumann served as an important catalyst for this interaction: he was a consultant to both the government and the computer industry, and he led the development of the computer at the Institute for Advanced Study at Princeton that served as a model for the development of a whole generation of computers, including the IBM 700 series computers and their successors, the ILLIAC, the Johnniac, the MANIAC, the ORDVAC, and the ORACLE. This highly beneficial interaction among government, universities, and industry continued through the decades of the 1950s and 1960s, but suffered a serious decline in the 1970s. Universities in particular have become isolated from the mainstream of supercomputer research and development. The leading supercomputer of the early 1970s was the Control Data Corporation 7600, but none of these computers was installed at an American university (although several were installed at foreign universities). Even today, only three American universities have on-site supercomputers. This unfortunate situation has resulted in fewer graduates with backgrounds in large scale scientific computing and fewer research projects being conducted in universities in this field than in previous decades. There is general agreement among these communities that it is essential to reestablish their close collaboration of earlier decades for the United States to maintain its leadership in supercomputer development, marketing, and applications.

III. Anticipated Requirements for Future Supercomputers

At the time this report was prepared, approximately 70 Class VI supercomputers have been sold in the U.S., the U.K., Germany, France, Sweden, and Japan. Only about 10% of these are at universities (three in the U.S.). Several reports have recently been written describing the need for more facilities in the academic environment (e.g., the Lax Report and the Press Report). Additional information on the applications in Government laboratories and industry is provided below.

A. Applications in Government Research and Development

The Federal Government purchased 29 CDC 7600 supercomputers, about 40% of the total number built. Manufacture of the CDC 7600 ceased in 1977, yet all 29 remain in productive use. The current generation of supercomputers includes the Cray-1 and the CDC Cyber 205. The Government has already acquired 29 of these, and several more are expected to be acquired over the next two years, including the recently announced Cray-XMP.

The Department of Energy is the largest user of supercomputers, with 13 Cray-1's and two Cyber 205's already installed. Nuclear weapons design and engineering are heavily dependent on the use of supercomputers, as are magnetic and inertial confinement fusion research and nuclear reactor development. NASA has two Cray-1's and one Cyber 205; they are used for atmospheric modeling, aeronautical design, and scientific research. The National Science Foundation has two Cray-1's installed at the National Center for Atmospheric Research, where they are used for atmospheric and other scientific research. The Department of Commerce uses two Cyber-205's, and a third will soon be added for NOAA's weather forecasting and oceanographic research. The Department of Defense uses a Cray-1 for weapons effects simulation and a Cyber 205 for weather forecasting.

The graph (Fig. 1) shows the present and projected Government installed base of supercomputers. DOE's designation of classes has been used to label each succeeding class of supercomputers. Class IV installations after the introduction of Class VI machines are not included because, by definition, these are no longer considered supercomputers. (In fact, there are many more installed Class VI and Class V machines, such as IBM 3081's and Cyber 176's than are shown in the graph).

Future projections are intentionally conservative on the graph. They are based largely on replacement of present machines following historical trend and new installations budgeted or planned. Not included is allowance for emerging applications and greater access, which could add another 10-20 machines by 1990. Planned new installations include NASA's Numerical Aerodynamic Simulation, DOE's Energy Research Network, NSF's supercomputer program, Commerce's National Bureau of Standards system, and DOD's Naval Research Laboratory.

Based on these projections, it appears that the direct Government requirement for supercomputers will be at least 30-40 machines from each new class. New applications could increase this to 50 Class VII computers.

ER52701T.01 83 Aug 11 W Thompson

Figure 1
GOVERNMENT INSTALLED SUPERCOMPUTERS

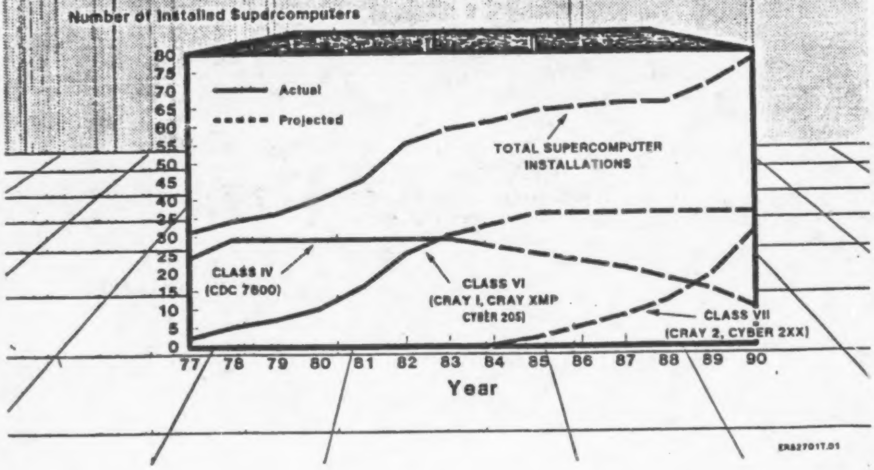

B. Commercial Applications of Supercomputers

Scientific applications have been developed and perfected by the national laboratories and universities in conjunction with research and development activities. The applications are gradually adopted by industry when the techniques are demonstrated to be valuable in modeling complex processes. Programs are typically made available to industry at a nominal cost through distribution organizations, such as DOE's National Energy Software Center, NASA's Cosmic, and DOC's National Technical Information Service.

Examples of current applications in use by industry are:

Aerospace Industry: Simulation of air flow around aircraft using finite difference and vortex tracking methods; structural analysis using finite element methods.

Petroleum Industry: Seismic exploration using inverse scattering techniques; tertiary recovery of oil from reservoirs using front tracking methods.

Automotive Industry: Combustion modeling for internal combustion engines using flame front tracking and chemical kinetics programs; computer aided design and testing of automobile parts.

Electronics Industry: Computer aided design of VLSI circuits; numerical solution of large sets of differential equations describing circuits using numerical linear algebra techniques.

Engineering and Manufacturing: Computer aided design of machine parts; testing structural stress-strain properties of structure using finite element methods; non-destructive evaluation using image reconstruction techniques.

Power Industry: Modeling of nuclear reactor safety systems; power grid optimization using linear programming techniques.

Chemical Industry: Process control and optimization; simulation of polymers; reaction diffusion modeling and turbulence.

Movie Industry: Computer graphics techniques for image generation and animation.

IV. Industry Outlook

A healthy U.S. supercomputer industry is necessary both for its effect on national defense issues and on the balance of trade. The computing equipment market, of which supercomputers are a bellwether, is one of the leading U.S. exports. Nuclear weapons design programs are highly dependent on advanced supercomputers and must utilize the highest capability available - this can only be assured if these are a U.S. product.

Supercomputers represent a small niche in the large computer market. The risks and costs associated with developing a new supercomputer system are large. Consequently, most large U.S. computer vendors have not entered the supercomputer field. An additional risk factor for the supercomputer vendors is their dependence for state of the art integrated circuits on merchant semiconductor manufacturers over whom they exert little, if any, influence. A similar situation exists for high speed peripherals.

In order to help remove some of the risk factor, there must be a well identified market - e.g., the Government sponsored market. If such a market is not clearly delineated, foreign manufacturers aided by their government's funding, and favorable export policies, will be the only companies willing to risk precious capital resources on poorly defined markets.

V. Goals

Because of the strategic importance of supercomputers to the long term strength of the United States, both economically and militarily, it is imperative that an adequate supply of domestically produced supercomputers be available in the U.S. Toward this end, the goal of U.S. Government policy should be:

To retain U.S. leadership in the development, production, and application of supercomputers.

The implementation of a policy based on this goal will be instrumental in:

1. Meeting U.S. government needs for advanced supercomputers and for the trained scientific and engineering personnel necessary for their use;
2. Meeting U.S. industry needs for advanced supercomputers for oil exploration, aircraft design, integrated circuit design, automobile engine design, etc.;
3. Maintaining this country's strength and vitality in supercomputer technology through support of basic research and education;

4. Providing incentives necessary to accelerate the development of advanced supercomputers by U.S. manufacturers; and

5. Retaining the U.S. lead in computer related technology as a whole through spin off from leading edge supercomputer technology.

VI. Recommendations

The Government has traditionally been a major purchaser of supercomputers and has supported hardware and software R&D that benefits supercomputer development. The intent of these recommendations are: (1) to focus Government activities sufficiently to provide incentives for accelerating supercomputer development, and (2) to focus Government-sponsored R&D in generic areas that are crucial for improved supercomputers but inadequately stimulated by market forces in the private sector. The development and production of complete supercomputer systems, including architecture, hardware, and software, cannot now be accomplished by any one organization. The best talents of U.S. industry, Government laboratories, and universities must be brought to bear in order to ensure expeditious development of complete, usable systems.

The recommendations are designed to address both the short-term and long-term requirements for accelerating supercomputer development as necessary to provide the large scale scientific computing needs of Government and to maintain U.S. leadership in this industry. The recommendations deal with several aspects, including manpower training, expansion of supercomputing applications, development of several 200x supercomputers by 1990, and long range R&D required for computers with capabilities beyond 200x. A 200x supercomputer refers to a system having a capability that is about 200 times that of current Class VI systems.

GOVERNMENT ROLE: Most of the recommendations provide recognition of the role that the government has played historically in supercomputer development rather than suggesting completely new government approaches. However, the recommendations do call for an increased emphasis on and an acceleration of this traditional government role in order to meet the future needs of the country and to retain U.S. leadership in this field.

1. Recommendation: Stimulate industrial development of several 200x scientific supercomputer systems by 1990.

The development of 200x supercomputers by 1990 is a challenging yet technically feasible goal. Several manufacturers have defined architectures that can be extended to this level of performance, but new development will be required in several areas: fast circuits with large scale integration, high density component packaging with adequate cooling, high performance peripherals, system integration of multiple processors, and improved software. Traditionally, supercomputer vendors have developed new architectures and have packaged components. Software development has been assisted by Government laboratories, with the vendors supplying only rudimentary software. Integrated circuits using off the shelf technology have been purchased from IC vendors. High performance peripherals have been developed and manufactured by vendors. The goal of the recommended program is to use the best capabilities of the supercomputer vendors, the IC vendors, peripheral vendors, Government laboratories, and universities.

The heart of this strategy is the guaranteed Government purchase or lease of three of the first supercomputers from each vendor who meets the 200x goal in the FY 1986-1990 time frame. The intent of this guarantee is to provide sufficient incentive for vendors to accelerate supercomputer development beyond current trends. The objective is to produce computer systems that not only meet Government requirements, but can also compete in the commercial market at a system price of $10-20 million.

The 200x goal is not well specified at this time and needs to be examined by an appropriate group of technical experts. Traditionally, new machines are "benchmarked" on a package of Fortran routines to be representative of the typical workload of existing large scale scientific computer centers. A recent suggestion is to judge supercomputer system performance on solving four or five scientific problems of national importance rather than using benchmark codes. The goal using this approach would be to execute the problem benchmark set at a speed-up of some 200 times that currently available, allowing for improvement in component speed, concurrent processing architectures and more efficient algorithms and languages. This new approach should be seriously considered.

In addition to the guaranteed acquisitions there are several other considerations in the recommended strategy.

a. Architecture and Packaging

The supercomputer vendors have the expertise to develop new architectures and to package components for the 200x systems, although this goal will provide challenges to the vendors in these areas. The guaranteed acquisition plus the remainder of the supercomputer market is considered to be an adequate incentive. No direct Government support for vendor R&D in these areas is recommended.

b. Integrated Circuit Development

Acquiring fast logic and memory circuits for advanced supercomputers has become a problem for U.S. vendors. IC manufacturers in the U.S. are not developing on a timely basis the specific technology required for supercomputers because this market is very small compared with the market for high volume products such as personal and business computers. Japanese IC manufacturers have put more effort into supercomputer technology; as a result, U.S. supercomputer vendors are already forced to obtain some of their memory IC's from Japan. This leaves the U.S. vulnerable to the trade policies of Japan, or of individual IC manufacturers who may refuse to export new IC technology until internal needs are satisfied. Importing logic chips is particularly dangerous, because the internal logic design of U.S. computers must be revealed in order to produce these chips.

Because IC development is so crucial to advancing supercomputer performance, and reliance on foreign vendors is dangerous, U.S. capability in this area must be examined carefully. At present, no U.S. supercomputer vendor is able to produce its own IC's. Cray Research has recently invested in setting up its own chip manufacturing facilities. However, high speed IC manufacturing is a most difficult undertaking, and many companies have failed in similar attempts. For these reasons, Government should support the base technologies associated with high speed IC development. R&D should be supported on silicon devices, advanced semiconductors such as gallium arsenide and new promising technologies. Cost to the Government could be kept to a minimum by piggybacking on existing Government sponsored or private activities. The component question should be examined in more depth.

c. High Performance Peripherals Development

Peripherals development suffers from the same problems as IC development. The market for high performance peripherals, such as disk storage, is relatively small. Consequently, they attract little development funding within the manufacturers. Recent experience in high performance peripherals development indicates that common supercomputer peripherals could be specified by the industry, resulting in a truly generic development program. Since advanced peripherals are essential to future supercomputer systems and market forces do not adequately push their development, some Government support for their timely development should be strongly considered. Two possible mechanisms are guaranteed government buys and government funding of research and development.

d. Software

Software requirements for future supercomputers include (1) improved operating systems and utilities to effectively use parallel processing, (2) an improved Fortran-like language with parallel capability, and (3) new algorithms to apply parallel processing techniques to single, large problems. Supercomputer vendors currently supply only rudimentary software; they receive considerable software assistance from Government laboratories. It is unlikely that vendors will be able to develop effective software for future supercomputers without continued assistance, yet software quality will be an important factor in international competition.

The strongest existing groups for developing supercomputer operating systems are in the national laboratories. Their expertise should be fully utilized. Alternatives should be explored to best use the resources of vendors, laboratories and unviersities to develop and maintain superior software.

A single common operating environment for supercomputers, including the operating system, utilities, and a Fortran-like language, would reduce the resources required by industry to develop software and would enhance competition by simplifying installation of supercomputers from multiple vendors at a single site. Discussions with supercomputer vendors indicate that, in principle, they welcome the idea of a common operating environment. The government should take a more assertive role leading to a common operating system and a high level user language.

Research on algorithms is required to enhance our ability to use parallel processing techniques in applications software. This research would benefit from the availability of experimental and prototype computers to researchers at government laboratories and at universities. Support of research of this type by the government should be accelerated. In the long term, vendors should be encouraged to produce high quality software.

e. Prototype Acquisition

Several vendors have announced supercomputers whose architecture may be extendible to 200x capability. These computers can legitimately be considered incremental steps leading to the 200x machine. Early procurement by the government should be assured to speed software development and provide feedback to the vendor as the 200x machines are being developed. The Government has traditionally been a "friendly buyer" of the first copies of supercomputers. This role is now even more important to encourage accelerated supercomputer development.

Purchase of prototype computers from domestic producers as soon as they become available will make increased capability available earlier and provide closer ties to the needs of mission oriented Government agency users. Their purchase will also provide incentive by providing partial early return on the investment for the development of hardware and software systems. In addition to purchases of prototypes, Government purchase of production versions of machines providing significant incremental advances will help maintain a financially healthy industry that can afford to invest in development of the 200x machine and even more advanced systems.

This type of advanced procurement notice will encourage industry research and risk-taking necessary to develop breakthrough units to achieve large increment increases in performance. Mission oriented agencies must work with computer developers to focus designs to meet critical government requirements in scientific, space, and military applications.

GOVERNMENT ROLES: (1) Provide support for supercomputer development through multiple *(at least three) procurements of the first 200x units* and through support of essential R&D in components, peripherals, and software; and (2) Provide support for supercomputer development through purchase of new prototype and production supercomputers which are significant steps toward the 200x goal.

2. Recommendation: Accelerate hardware and software development through increased Government support of experimental supercomputer systems.

The purpose of the program in experimental computers is to provide an early opportunity for government laboratories and universities to explore new architectures and advanced concepts for a range of applications or to accomplish a unique task specific to the user organization. Experimental computers are defined as systems with hardware that is thoroughly debugged and warranted by the vendor, but with operating and support software that is not mature and in some cases may be totally lacking. Such systems are procured on a research and development basis rather than in response to program workload requirements, and will be used to support laboratory and university research on architecture, algorithms, and develoment of software rather than to solve production workload problems. Experimental computers may be a precursor to future production systems, but may differ in terms of engineering and performance from later production versions.

Immediate experience with experimental processors is required to determine how to use new architectures and to guide specification and evaluation of new systems such as the 200x supercomputers. Several machines have been developed recently which, although not supercomputers, have very interesting architectures and present opportunities to learn about parallelism concepts in particular. Several of these systems should be installed in government laboratories and universities and made available, through computer networks to researchers for exploration of new algorithms, languanges, software tools, and applications.

Care should be taken to assure that a clear distinction is made between the acquisition and use of research and experimental systems versus the acquisition and use of prototype and production computers, in order to avoid misunderstandings concerning the expectation of results.

GOVERNMENT ROLE: Provide increased resources for and access to advanced research and experimental supercomputer systems.

Increased funding is required for experimental system developments and for networks for access to these systems by the Government, university, and computer industry R&D community.

3. Recommendation: Provide improved access for expansion of applications and manpower training.

At the present time, very few universities have access to supercomputers. Because of this lack of access, there is very little training of university students in the use of supercomputers. Consequently, government laboratories and industry have to expand substantial time, one to two years, training new hired scientists and engineers to use large scientific computers. Further, the lack of access of university researchers to supercomputers means that a large number of scientists do not utilize one of the most powerful scientific tools available today. In contrast, some European countries have placed new supercomputers in their universities, and U.S. researchers are complaining about their inability to compete in forefront research with their better equipped European colleagues.

Providing universities access to supercomputers as soon as possible is of paramount importance to the country's R&D activities as well as to manpower training. In addition, the availability of supercomputers to the many bright, young students and faculty would undoubtedly result in unanticipated applications for scientific computing. These new applications could broaden the market, increase supercomputer sales in the future, and strengthen the ability of U.S. industry to compete in worldwide markets.

GOVERNMENT ROLE: The mechanism by which access should be provided is the subject of an interagency committee and will not be addressed in detail here. However, the access question is considered to be of the utmost importance to the health of the supercomputer industry.

4. Recommendation: Enhance long range basic R&D for scientific supercomputing.

Supercomputer architecture ideas for the 1990's must come from research and development activities of the 1980's. Currently there are several interesting architectural studies being pursued in universities and in Government and industry laboratories. The areas under investigation include problem decomposition techniques, algorithms, languages, and software environments for massively parallel computer architectures. These new architectures, employing hundreds or thousands of fast processors working concurrently to solve a single problem, require a thorough reconsideration of methods for the decomposition of problems into parallel operations, the creation of numerically stable algorithms for solving the problem components, the development of new software tools for ensuring the robustness and correctness of the implementation, and the development of new techniques for representing the solutions.

The R&D projects envisioned include interdisciplinary teams of computational scientists and computer scientists working on all aspects of large scale scientific computing problems. The design and construction of several potentially strong candidate experimental machines will be an important proof of concept activity.

University researchers will play the major role in generating ideas and experimental software and in training graduate students in generating new architectural concepts and applications. Government laboratory staff are in the forefront of tackling real world, large scale scientific problems and have unique resources for participating in these research projects. Industry likewise has a unique role in providing state of the art production and testing facilities and would stand to reap great benefits in understanding future architecture and software issues that tend to limit industry use of supercomputers currently. The transfer of technology from the academic and laboratory research environment to industry would be as rapid as possible through these cooperative projects.

GOVERNMENT ROLE: Government agencies responsible for supporting basic research should expand significantly their support of computational aspects of research programs in the physical, mathematical and social sciences. In particular, the present research programs in computational mathematics, algorithms, software engineering, and the development of advanced experimental machines based on novel architectural ideas should be enhanced to provide the future advances in computational science and engineering. Both the research aspects and the training of graduate students and postdoctoral researchers should be emphasized.

This program will be coordinated by the participating agencies to share common facilities where possible, such as VLSI design and fabrication facilities, experimental machines, and computer networks.

5. Recommendation: Research and development tax incentives for industry involved in the development of supercomputer systems components should be explored by an appropriate group, e.g., a committee with representatives from OMB, the Department of Commerce and the Department of Justice.

6. Recommendation: Some process whereby the acceleration of export licensing can be achieved should be explored by an appropriate group so that U.S. vendors are not unduly penalized in competing for foreign customers.

7. Recommendation: A permanent interagency group should be established to coordinate individual agency activities in supercomputing as necessary to implement the recommendations in this report.

All Federal agencies directly involved in supercomputing R&D should be represented in the coordinating group, which will report directly to FCCSET. This group should function at both the policy and the technical level, and act to coordinate procurement of new experimental supercomputers, software development, networking and access, and long range R&D.

FCCSET Working Group

James F Decker

Dr. James F. Decker, Chairman
Department of Energy

Dr. Robert E. Kahn
Department of Defense

E. F.

Dr. E. F. Infante
National Science Foundation

Leonard

Dr. Leonard A. Harris
National Aeronautics & Space
Administration

Dr. Joe Smagorinsky
Department of Commerce

George W

Mr. George Deskin
Director, Planning & Policy
Staff
Intelligence Community Staff

Appendix A

The Japanese have begun a major effort to become the world leader in supercomputer technology, marketing, and applications. This effort includes not only the development of supercomputers by the three largest Japanese computer firms but the support by the Ministry of International Trade and Industry (MITI) of three national projects affecting supercomputer technology.

1. NATIONAL PROJECTS

1.1 The Super-Speed Computer System (SSCS)

Six major Japanese computer firms (Fujitsu, Hitachi, NEC, Mitsubishi, Oki Electric, and Toshiba) are collaborating under the leadership of the Electro-technical Laboratory (ETL) to develop the technology that will make it possible to build a supercomputer capable of performing at a peak rate of 10 billion floating-point operations per second (FLOPS). This project is supported by about $100 million from MITI and is to be completed in 1989. If successful, this project would lead to the development of computers some 100 times as powerful as the Cray-1. The six companies involved in this project have revenues that are about 8 times as large as the revenues of the American supercomputer vendors; with the collaboration of a national laboratory (ETL) and the financial backing of MITI, the resources backing this project can be estimated to be at least 10 times as great as those available to the American supercomputer vendors.

1.2 The Fifth-Generation Computer System (FGCS)

Eight organizations (Fujitsu, Hitachi, NEC, Mitsubishi, Oki Electric, Toshiba, ETL, and Nippon Telephone and Telegraph (NTT)) are collaborating in the research and development of technology intended to revolutionize the design of computers using artificial-intelligence concepts. A new institute, the Institute for New-General Computer Technology (ICOT) has been chartered to oversee this project. Total funding for the project has not been specified, but the funding for the first five years of this ten-year project will be about $430 million.

1.3 New-generation technologies project

This project includes much more than is of interest for supercomputers, but it includes the development of high-speed components, three-dimensional device geometries, and components able to withstand hostile environments. The high-speed components for the supercomputer project include not only extensions of silicon technology but efforts to develop high-speed gallium arsenide (GaAs), high-electron-mobility transistors (HEMT), and Josephson Junction technology.

2. COMMERCIAL PROJECTS

2.1 Fujitsu

Fujitsu has announced the VP-200 and VP-100 computers that will be available for delivery in the fourth quarter of 1983 in Japan; marketing of these computers will occur outside of Japan in 1984. These computers have peak-performance rates of about 500 and 250 million FLOPS, respectively; this compares to about 400 million FLOPS for American supercomputers. These computers will execute the IBM instruction set.

2.2 Hitachi

Hitachi has announced the S810-20 and S810-10 computers that will be available for delivery in the fourth quarter of 1983 in Japan and in other nations in 1984. These computers have peak-performance rates of about 600 and 300 million FLOPS, respectively, and will execute the IBM instruction set.

2.3 Nippon Electric (NEC)

NEC has announced that they will deliver two supercomputers in the first quarter of 1985: the SX-2, and the SX-1, with peak performances of 1300 and 650 million FLOPS, respectively. These computers are not IBM-compatible.

3. SITUATION AUDIT

3.1 Japanese strengths

In-house semiconductor capability. Hitachi, NEC, and Fujitsu rank 3rd, 4th, and 8th in the volume of semiconductors manufactured in the world. This gives these companies an in-house components development capability that can respond quickly to the needs of their systems designers for high-speed logic and memory components, and it will allow these companies to shorten the design cycle for new generations of supercomputers compared to their American competitors. This is a capability lacked by the American supercomputer vendors and is a key reason why American supercomputers have evolved so slowly in the last decade. The lack of an in-house high-speed components capability is the Achilles heel of the American supercomputer effort. Cray research is attempting to develop such a capability, but to expect this small company (1/10th the size of Hitachi, for example) to compete with some of the world's leading merchant semiconductor houses is unreasonable.

IBM-compatibility. Two of the three announced Japanese supercomputers are capable of running IBM-compatible software. This is an option available to supercomputer customers that has hitherto been lacking. This will meet the needs of organizations where the computer centers use IBM-compatible systems and also need supercomputer speeds in some of their programs. There is an estimated $200 billion worth of IBM-compatible software available in the world and this will make a software-rich supercomputer environment available to supercomputer users for the first time. It is entirely possible that a significant fraction of the forecasted growth in supercomputer usage could be in the IBM-compatible area. Because no American company offers an IBM-compatible supercomputer, this market will go by default to the Japanese unless some action is taken by American vendors to correct this deficiency in their project offerings.

Nationally subsidized supercomputer R&D. The risks associated with both the Super-Speed Computer System and the Fifth-Generation Computer System are being shared by the Government of Japan and the Japanese computer industry in the form of Government subsidies. In the long term, it is unlikely that the relatively small American supercomputer industry can compete successfully against an effort having the backing of the Government of Japan and an order of magnitude more resources than are available to the American supercomputer vendors.

3.2 Japanese weaknesses

Lack of experience with supercomputers. The Japanese have two American supercomputers installed and two of the supercomputer companies have previously developed vendor processors, so they are not without experience in the field. However, compared to the long history of design, development, and application of supercomputers in the U.S., the Japanese are initially at a decided disadvantage. However, the ability of the Japanese to do "reverse engineering" on competing products has been demonstrated many times, and the Japanese relative lack of experience will be a short-lived advantage for the Americans.

Lack of an innovative tradition. The Japanese have moved quickly into the supercomputer market by following the lead of the American designers and attempting to develop a superior product on the basis of the American designs. However, the next generation of supercomputers will require true innovation in the development of parallel processors, for which the Japanese cannot simply copy American designs. It remains to be seen how well the Japanese meet this challenge.

Weak higher education system. Compared to the American and European systems of higher education, the Japanese universities are decidedly weak in computer science education. The Japanese universities graduate more engineers than do the Americans, but the quality of computer-science education is higher in the United States than in Japan. This is an advantage that clearly must be exploited in American planning.

Lack of a supercomputer marketing base. The Japanese do not have a supercomputer marketing base, i.e., an infrastructure of sales, consulting, and maintenance personnel. A supercomputer by itself is of no use: it must have extensive support, and the Japanese have yet to develop that support system.

3.3 Japanese opportunities

Minimal competition. The supercomputer industry differs from most areas of the computer industry in that there is a small market and only minimal competiton for that market. This makes the American supercomputer vendors vulnerable to a concerted effort by a foreign vendor to marshall larger resources than are available to the American vendors and overwhelm them in the long term.

IBM-compatible supercomputer market. The last IBM-compatible supercomputer-class machine was the IBM 370/195, marketed in the early 1970s. Since then, the customer who needed supercomputer speed bought that speed at the expense of a software-poor computing environment, because the market forces were not strong enough to pay for the development of extensive supercomputer software. IBM has ignored this market, apparently for two reasons: (1) it was only a small market, and (2) it was contrary to IBM's and the industry's interest for the supercomputer vendors to be overwhelmed by the IBM giant. Further, IBM customers, because they were willing to choose the high-end IBM machines (such as the 308X) in order to retain the advantages of IBM-compatible software. With the entry of the Japanese supercomputer offering, the customer can now have both supercomputer speeds and software richness.

Internal Japanese requirements. A part of the Japanese motivation for developing their own supercomputers is to meet the needs of their own users. If ten Japanese customers bought American supercomputers, then some $100 million would flow out of Japan. Spending $100 million to develop their own competing products would save that investment for Japanese companies. Further, Japan is not just interested in the supercomputer market as such; rather they are interested in supercomputer applications, including aerospace, energy systems design, weather forecasting, integrated circuit design, fluid dynamics, structural design, reactor safety, and exploration for natural resources.

4. RESPONSES

In addition to the responses generated in the U.S. by these Japanese actions, there have been responses in Western Europe.

4.1 European Economic Community (EEC)

The EEC is completing plans to begin a program called "ESPRIT" (European Strategic Plan for Research in Information Technology) that calls for collaboration in "pre-competitive" R&D, very much in the Japanese style. This project is to run five years and be funded at about $1.5 billion, with about half of the funding to come from the governments in the EEC and half from the participating companies. This project is to begin in January 1984.

4.2 Great Britain

Great Britain has undertaken a project in response to the Japanese fifth-generation project, to be funded at about $550 million over five years, with the funding costs to be shared by the British Government and private companies. This project will focus on the following: (a) software engineering, (b) man-machine interface, (c) artificial intelligence, and (d) VLSI. Britain will also participate in the ESPRIT project.

4.3 France

France has passed a "Science Program Law" that calls for increasing the level of national support from 1.8% of GNP to 2.5% of GNP by 1985 (this is about the level of expenditures in the U.S. and other western nations, and Japan). This funding will be used to increase funding and staffing at national research centers and for grants, loans, and tax exemptions for private industry. France is considering building their own supercomputers, but no decision on this has been announced yet. France will participate in the ESPRIT program.

4.4 Germany

In addition to participating in the ESPRIT program, Germany has recently increased its support for parallel processing by $4 million annually.

4.5 Other

Both Korea and Taiwan have announced plans for government support of collaborative R&D programs.

5. SUMMARY

The scope of the Japanese supercomputer effort is much larger than the scope of the American supercomputer effort. This larger set of resources, if applied over the long term, will simply overwhelm the relatively small American supercomputer effort that currently exists unless there is some appropriate response from the American supercomputer industry, the government, and academia.

ACADEMIC SUPERCOMPUTER R & D PROSPECTS

for

Committee on Science and Technology
U. S. House of Representatives

by

David Kuck

Laboratory for Advanced Supercomputers
Department of Computer Science
University of Illinois at Urbana-Champaign
Urbana, Illinois 61801

October 19, 1983

ACADEMIC SUPERCOMPUTER R & D PROSPECTS .

Supercomputer research in U. S. universities must be supported for several reasons: first, to train users and designers of supercomputer systems, secondly, to provide objective analysis of the design and performance of commercially available architectures and software, and finally, to provide designs for and prototypes of future generation architectures and software.

The number of universities that can effectively engage in each of these areas decreases as we proceed from points one to three. Thus, there are dozens of universities that can train supercomputer users, and even designers, given adequate access to supercomputer systems. Fewer can mount serious analysis projects, and perhaps only half a dozen are presently capable of designing and building working prototype systems that would advance science or the state of the art.

Several university groups have been seriously considering major prototyping efforts for the past twelve to eighteen months. The likelihood of federal government funds for such efforts has alternated from being seemingly assured to appearing hopeless over much of that time. However, the excitement and drive of a major research group cannot be maintained indefinitely with periodic hopes of future funding. Unless it becomes clear in the next few months that major new initiatives will occur in FY 1985, most of these groups will lose their momentum (including key staff members, etc.). Rebuilding such groups would take years.

The total funding needed is not large by federal government standards. A typical prototyping project needs about ten million dollars over a five-year period. Major software and algorithm projects are less expensive. Thus, six prototype supercomputers, together with a number of major software projects, as well as analyses of existing machines could be supported for twenty million dollars per year. Such an effort would be very effective in helping the U. S. maintain world leadership in supercomputing in the 1990's.

Academic supercomputer researchers have every reason to be worried about the current budget situation. At the August meeting of the National Science Board, it was decided to drop new NSF plans to fund supercomputer research at the prototype level. The NSB only approved plans to buy and provide access to commercially available systems. The stated reason for this decision was that DARPA would fund research and development efforts for new supercomputers.

This appears to be quite far from the truth of the matter. DARPA's position seems to be that architectures for symbolic computation and artificial intelligence, plus signal processing, are the key considerations. Dr. Robert Cooper of DARPA has stated that scientific supercomputing is well understood, but that these other areas need much work. If that were true, then the Cray and CDC supercomputers would be delivering one or two hundred megaflops each, through Fortran compiled code, instead of about ten megaflops. Furthermore, the key reasons that these machines perform slowly on many scientific and engineering problems are essentially the same reasons that prevent them from doing well on artificial intelligence problems. Thus, many of the ideas that are currently under active study by the university computer design groups referred to earlier will have wide applicability in future computer architectures for scientific and engineering applications as well as symbolic computation (although they appear ineligible for DARPA funds).

The Department of Energy has recently presented plans for funding new supercomputer research to the OMB. This would provide a major expansion of their current modest research support in high speed computation. Based on the current situations at NSF and DARPA, and on the long history of DOE support of supercomputer development and use (stemming from AEC times), the DOE seems to be a very good candidate for a major new activity in supercomputer research and development. Through the national labs, the DOE has served as a focal point in recent years by organizing meetings of key researchers, providing benchmark programs and computer time, and encouraging new architectures generally. Thus, DOE funding seems not only to be the last hope for innovative supercomputers, but based both on distant and recent history, a very propitious choice.

If academic supercomputer architecture research is not expanded to the prototyping phase, several dire consequences are likely. First, the student training function will not be able to expand to meet the rapidly increasing demands. That these demands will continue to grow can be seen, for example, by observing the explosion of array processor production (at least forty companies produce them today). Today's array processor users become tomorrow's supercomputer users. Secondly, the important university groups will either collapse or leave their universities to attempt private commercial activities. While this could lead to some interesting new products, it would have a crippling effect on the production of students, faculty consulting with industry, and the flow of ideas into the public domain.

New funds for software and applications program development for supercomputers are also urgently needed. The lack of software for exploiting these systems currently provides a serious barrier to entry for new users. It has been demonstrated that many of these software problems can be solved. What is needed is more work on implementing and extending the known solutions. Traditionally, vendors have provided new supercomputer hardware with very weak software. The time required to implement good software has been about equal to the time for a new supercomputer to be developed. Thus, software has lagged about one generation behind hardware. Breaking out of this vicious cycle will allow many more scientists and engineers to exploit supercomputers.

We need a broad program, announced soon and funded in FY 1985, with some starting funds in FY 1984, to produce prototype supercomputers and software for scientific and engineering applications. The funds should be available to university researchers for work alone or in conjunction with established industrial organizations.

RESOURCE MATERIAL TO

COMMITTEE ON SCIENCE AND TECHNOLOGY
HEARINGS ON SUBJECTS RELATING TO
UNIVERSITY SUPERCOMPUTING

OCTOBER 26 & 27, 1983

Robert W. Hayman
Assistant Vice President
for Computing Resources

Colorado State University
October 21, 1983

This paper is addressed to the general indictment that university computer centers are unable to serve a sophisticated university research community. These charges have been raised over the past several months in connection with the need to provide supercomputing into the university community of scientists and engineers. Recently, specific federal initiatives are preparing to provide supercomputer access to the university community. A legitimate question is raised concerning the funding of facilities to provide this service. Should the facilities be operated by federal laboratories who historically operate supercomputer facilities or should the facilities be operated by the universities themselves? Thoughtful consideration of this question raises strong reservations about the ability of the university computer centers to operate the supercomputer facilities and to manage the distribution of services to a geographically dispersed community. The comments to follow here will address this issue.

Except for a brief period during the mid-1960's, university computer centers have been in a poor position to provide adequate support to the campus research programs. To those of us who have been involved with this problem over the years, the reasons for this seem clear; however, these reasons need to be stated in terms which are just as clear for others to appreciate.

Three primary conditions exist to create major difficulty. First, university computer centers have been seriously impaired by continued underfunding. Second, with minor exceptions, no university computer center has been allowed to organize itself with a primary objective to support the local research community. Third, laboring under the first two constraints, a process of allocation of available resources has been required to distribute available resources among competing campus interests. The allocation process, in general, has been very poorly managed.

The condition of underfunding of campus computing facilities will not be belabored here since the fact has been readily recognized and documented over a period of time. The most important point to be observed in this connection is that the underfunded condition has affected the research community disproportionately. The reason for this is associated with the source of budget support for the computing centers. Such support is offered primarily to accomplish the administrative mission of the university and to provide traditional classroom computing. In fact, there are many states who simply refuse to provide any visible budget to support a serious research computing function. The state governments appear to believe that research computing should be funded from the same sources as the research programs themselves. Under these conditions, research computing is supported as an accommodation rather than a recognized mission. It is astonishing to learn that the judgement or managerial ability of the campus computing center director is thought to be a significant basis for these conditions.

Given the stated attitude of the state governments and the philosophy of the federal funding agencies, there should be little mystery with the observation that campus computing has been unable to organize a center which emphasizes the services required by the campus researchers.

Digressing in time to the mid-1960's, we can observe that the converse of the present-day conditions existed for a time. During this period the NSF computing initiatives were aimed more specifically at the research community with state funds being added to campus computing budgets as a secondary support. Consequently, the research community 'ruled the roost' and classroom and administrative computing functions were considered secondary. During this 'golden period' campus computing personnel demonstrated very clearly their ability to service the research interests and to master the computing systems which they had been provided. The vendor's operating systems were not considered sacred and were attacked routinely. System integration problems were attacked forcefully, data communication technologies were advanced and the foundation for important software libraries was laid. A few universities ventured in to the domain of high performance computing systems and, as well as any organization, proved their ability to comprehend these systems and to advance their technology.

Returning to more recent times, we should not wonder at the difficulty in attempting to manage the overstressed campus computing facility or the difficulty in allocating available computing resources among the three functional areas of classroom, research laboratory and administration. It should not be had to understand that very few computer center directors have been up to this task. This observation, which is widely made, should be corrected to state that very few universities have been up to the task. Somehow, the judgement and credibility of computer center management has been made victim.

The problems associated with management of the scarce computing resources has been greatly emphasized by federal fiscal policy which is required that dollar values be put to computer services and equally applied in each of the three campus computing arenas (Circular A-21, Article J-38). On the one hand, the state general fund budgets have provided little support for research computing, neither would federal funding agencies on the other hand and the federal fiscal policy has allowed no maneuver room or solution space which could advantage the campus researcher.

Given these conditions, the frustration of the campus researcher has continuously increased to the present time and feeds the current trend toward computing self-sufficiency on the part of the individual research program. It must be noted that the same historical conditions which support the current trend will also prevent individual research programs from sharing computing resources among themselves. Furthermore, the nature of the campus research will simply reduce to the capability of the tools which the researcher has at his disposal. Therefore, interest in computing self-sufficiency has the effect of reducing the scope of research, scientific inquiry and education to a restricted set of affairs.

The federal supercomputing initiatives under consideration at the present time are advocated around the principle that the university research community should resume a major role in the supercomputing technologies. In other words, it has been recognized that a range of computing tools should be available on a routine basis to the university scientists and engineers.

It is certain that several universities will attempt to participate in the various initiatives to support large scale computing for the university environment. At least three universities are preparing to do so at the present time. These efforts suffer under the general indictment identified at the outset of this paper. I respectfully submit that the adverse conditions discussed herein are inclined to precondition consideration of the university supercomputing proposals in a negative sense. All proposals should be considered on the basis of carefully constructed criteria and evaluated accordingly. In a general sense, criteria should be composed to reveal the following program attributes:

- adequate understanding of the mission
- realistic program design
- sufficient latitude to perform the program

From these criteria, reasonable questions can be composed and objective inquiry performed which will identify the ability of a given site to operate a quality supercomputing program.

Robert W. Hayman
October 21, 1983

○